No. 2023-1266

In the
United States Court of Appeals
for the Federal Circuit

VLSI TECHNOLOGY LLC,

Appellant,

v.

INTEL CORPORATON,

Appellee.

Appeal from the United States Patent and Trademark Office,
Patent Trial and Appeal Board of 2020-1744, -1745 (Remand),
Proceeding Nos. IPR2018-01312 and IPR2018-01661.
The Honorable Bart Alex Gerstenblith, Kimberly McGraw and Lynne Pettigrew,
Administrative Patent Judges.

BRIEF OF APPELLANT
VLSI TECHNOLOGY LLC

KAMRAN VAKILI
(kvakili@irell.com)
IRELL & MANELLA LLP
840 Newport Center Drive
Suite 400
Newport Beach, California 92660
(949) 760-0991

Attorney for Appellant
VLSI Technology LLC

PATENT CLAIMS AT ISSUE

This appeal concerns claims 1-5, 12-16, 18, and 20 of U.S. Patent No. 8,020,014. These claims are reproduced below:

1. A method for power reduction, the method comprises:

   selectively providing power to at least a portion of a component of an integrated circuit during a low power mode; and,

   determining whether to power down the at least portion of he component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode.

2. The method according to claim 1 wherein the component is a cache memory.

3. The method according to claim 2 wherein the component is a cache memory that comprises independently powered portions.

4. The method according to claim 2, wherein the determining is responsive to an amount of information that is stored only in the cache memory.

5. The method according to claim 1, wherein the determining is responsive to an amount of information associated with dirty bits.

12. A device having power reduction capabilities, the device comprises:

   power switching circuitry adapted to selectively provide power to at least a portion of a component of the device during a low power mode, having power management circuitry adapted to determine whether to power down at least the portion of the component during a low power mode in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode.

13. The device according to claim 12 wherein the component is a cache memory.

14. The device according to claim 13 wherein the component is a cache memory that comprises independently powered portions.

15. The device according to claim 13, wherein the determination is responsive to an amount of information that is stored only in the cache memory.

16. The device according to claim 12, wherein the determination is responsive to an amount of information associated with dirty bits.

18. The device according to claim 12, further adapted to generate cache statistics, and wherein the determination is responsive to the generated cache statistics.

20. The device according to claim 12 further adapted to monitor integrated circuit behavior, and wherein the determination is responsive to results of the monitoring.

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

CERTIFICATE OF INTEREST

Case Number 2023-1266

Short Case Caption VLSI Technology LLC v. Intel Corporation

Filing Party/Entity Appellant, VLSI Technology LLC

Instructions: Complete each section of the form. In answering items 2 and 3, be specific as to which represented entities the answers apply; lack of specificity may result in non-compliance. **Please enter only one item per box; attach additional pages as needed and check the relevant box**. Counsel must immediately file an amended Certificate of Interest if information changes. Fed. Cir. R. 47.4(b).

I certify the following information and any attached sheets are accurate and complete to the best of my knowledge.

Date: 01/03/2023

Signature: /s/ Kamran Vakili

Name: Kamran Vakili

1. Represented Entities. Fed. Cir. R. 47.4(a)(1).	**2. Real Party in Interest.** Fed. Cir. R. 47.4(a)(2).	**3. Parent Corporations and Stockholders.** Fed. Cir. R. 47.4(a)(3).
Provide the full names of all entities represented by undersigned counsel in this case.	Provide the full names of all real parties in interest for the entities. Do not list the real parties if they are the same as the entities. ☐ None/Not Applicable	Provide the full names of all parent corporations for the entities and all publicly held companies that own 10% or more stock in the entities. ☐ None/Not Applicable
VLSI Technology LLC	VLSI Technology LLC	CF VLSI Holding LLC
	CF VLSI Holding LLC	

☐ Additional pages attached

4. Legal Representatives. List all law firms, partners, and associates that (a) appeared for the entities in the originating court or agency or (b) are expected to appear in this court for the entities. Do not include those who have already entered an appearance in this court. Fed. Cir. R. 47.4(a)(4).

☐ None/Not Applicable ☐ Additional pages attached

Benjamin Hattenbach, Irell & Manella LLP	Amy E. Proctor, Irell & Manella LLP	Michael R. Fleming, Irell & Manella LLP

5. Related Cases. Provide the case titles and numbers of any case known to be pending in this court or any other court or agency that will directly affect or be directly affected by this court's decision in the pending appeal. Do not include the originating case number(s) for this case. Fed. Cir. R. 47.4(a)(5). See also Fed. Cir. R. 47.5(b).

☐ None/Not Applicable ☐ Additional pages attached

VLSI Technology LLC v. Intel Corporation, Case No. 5:17-cv-05671 (N.D. Cal., Oct. 2, 2017		

6. Organizational Victims and Bankruptcy Cases. Provide any information required under Fed. R. App. P. 26.1(b) (organizational victims in criminal cases) and 26.1(c) (bankruptcy case debtors and trustees). Fed. Cir. R. 47.4(a)(6).

☑ None/Not Applicable ☐ Additional pages attached

TABLE OF CONTENTS

Page(s)

TABLE OF AUTHORITIES

Page(s)

Cases

Statutes

Rules

STATEMENT OF RELATED CASES

A prior appeal from the same *inter partes* reviews ("IPRs") was previously before this Court (consolidated appeal nos. 2020-1744, -1745; decided May 26, 2021; before Judges Newman, Lourie, and Dyk; *Intel Corp. v. VLSI Tech. LLC*, 858 F. App'x 349, 352-355 (Fed. Cir. 2021)). No other appeal was previously before any other appellate court. U.S. Patent No. 8,020,014 ("the '014 patent") was previously asserted in *VLSI Technology LLC v. Intel Corporation*, No. 5:17-cv-05671 (N.D. Cal.), having been dismissed from the case by an order dated April 18, 2023. Counsel is not aware of any other cases that could directly affect or be directly affected by the decision in this appeal.

JURISDICTIONAL STATEMENT

This is an appeal from the Final Written Decisions on remand issued by the Board in two IPR proceedings involving U.S. Patent No. 8,020,014 (IPR2018-01312 and IPR2018-01661). The Board issued its final written decisions on October 11, 2022 pursuant to 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73. Appx1-59. VLSI Technology timely filed its notices of appeal on December 8, 2022. Appx626-628, Appx2558-2560; *see also* 37 C.F.R. § 90.3(a)(1). This Court has jurisdiction under 28 U.S.C. § 1295(a)(4)(A) and 35 U.S.C. §§ 141(c), 319.

INTRODUCTION

In the IPR proceedings below, Appellant Intel Corp. ("Intel") challenged claims 1-5, 12-16, 18, and 20 of Appellee VLSI Technology LLC's ("VLSI") U.S. patent 8,020,014 ("the '014 patent") based on references teaching approaches to power savings in cache memory systems that differed greatly from each other and from the '014 patent. In its original Final Written Decision, the Board properly ruled against Intel's allegation that one reference, Takahashi (U.S. patent no. 5,761,715; Appx869), rendered the challenged claims obvious, rejecting Intel's misinterpretation of Takahashi's "predetermined value" as a proxy for the claimed "estimated power gain." The Board also found that Intel's second primary reference, Hu ("Let Caches Decay: Reducing Leakage Energy via Exploitation of Cache Generational Behavior," Appx883), did not teach using a quantity called the *L2Access:leak* ratio to determine whether to power down portions of a cache. On appeal, this Court affirmed the Board's finding regarding obviousness over Takahashi alone, but reversed the Board regarding Hu teaching to use *L2Access:leak* for determining whether to power down. Accordingly, the case was remanded for the Board to determine whether the combination of Takahashi and Hu taught all of the claim elements, and whether a person of ordinary skill would have

had a motivation to combine the references and a reasonable expectation of success in doing so.

On remand, the Board reversed its prior position, finding that the combination of Takahashi and Hu taught all elements of the challenged claims, and that a person of ordinary skill would have been motivated to make the combination with a reasonable expectation of success. That decision was in error and unsupported by substantial evidence in at least four significant respects.

First, the Board relied on Hu's *L2Access* and *leak* quantities for the claimed "an estimated power gain and an estimated power loss resulting from powering down." But these are intrinsic quantities, unaffected by the operational factors "resulting from powering down" that would be essential to determine power gain and loss. It is undisputed that *L2Access* is the energy consumed *per access* to the L2 cache, and that *leak* is the leakage energy *per clock cycle*. Without information about how many L2 accesses would be required upon shutting down or about how long (*i.e.*, how many cycles) the shutdown would last, these quantities are simply incapable of reflecting power gain or power loss resulting from powering down portions of the cache. This critical shortcoming in the evidence relied upon by the Board fundamentally undermines its ultimate conclusion of obviousness.

Second, the Board relied on specific numerical values for *L2Access* and *leak*, calculated by Hu based on a literature review, as the "estimated" values alleged to correspond to the claimed "estimated power gain" and "estimated power loss." However, it did not rely on these for the same claim limitations in the context of the claims as wholes, instead relying on entirely different values that Hu simply made up for its simulations (the Board previously having acknowledged these to be "assumed" values in its prior Final Written Decision). This approach, dissecting the claim and relying on disjointed teachings of the cited art, reveals that the Board's ultimate conclusion of obviousness is not supported by substantial evidence.

Third, the Board credited Petitioner's argument that alleged similarities between Takahashi and Hu support a conclusion that a person of ordinary skill would be motivated to combine the references. However, the Board's conclusion was premised on two demonstrably incorrect findings: (1) that Takahashi did not teach powering down a cache way before measuring the cache-miss rate, and (2) that Takahashi's stated purpose did not include avoiding increasing the cache-miss rate. A correct disposition on these issues establishes that the proposed combination of references would fundamentally alter Takahashi's principle of operation and render it unsuitable for its stated

purpose, demonstrating that there would be no motivation to make the combination.

Fourth, the Board failed to provide adequate analysis supported by substantial evidence to support its conclusion that a person of ordinary skill would have had a reasonable expectation of success in combining the proposed references to reach the claimed invention. Petitioner largely neglected to even allege, let alone prove, a reasonable expectation of success in its IPR briefing, and the Board's findings on this issue are either entirely conclusory, or rely on Petitioner's deficient showing and allegations of similarities between the references based on grossly overgeneralized characterizations.

Accordingly, for each of these four independent reasons, this Court should reverse or vacate the Board's decisions.

STATEMENT OF ISSUES ON APPEAL

1. The Board relied on Hu's *L2Access* and *leak* for the claimed "an estimated power gain and an estimated power loss resulting from powering down." Did the Board err in relying on quantities that lack any information about operational conditions, rendering them incapable of being estimated power gain and loss "resulting from powering down"?

2. The Board relied on certain calculated values of Hu's *L2Access*, *leak*, and the *L2Access:leak* ratio for the claim elements "estimated power gain"

and "estimated power loss" in isolation, but relied on different, unrelated values of *L2Access:leak* for the same claim elements in the context of the claim as a whole. Did the Board err in dissecting the claims and relying on disconnected teachings in Hu for the same claim limitations, where Hu never shows the calculated values of *L2Access:leak* being used in connection with any determination of whether to power down?

3. In finding motivation to combine Takahashi and Hu, the Board discounted Patent Owner's argument that the combination would violate Takahashi's basic principle of operation and frustrate its ability to achieve its stated purpose. Did the Board err in reaching this conclusion based on its incorrect factual findings denying (1) that Takahashi teaches powering down a cache way before measuring the change of cache-miss rate; and (2) that Takahashi's stated purpose is to avoid increasing the cache-miss rate?

4. The Board found that a person of skill in the art would have a reasonable expectation of success in combining the cited art to reach the claimed invention. Did the Board err in reaching this conclusion where Petitioner failed to allege or prove such reasonable expectation, and the Board either neglected to provide supporting analysis and evidence, or provided only conclusory assertions based on gross overgeneralizations of the references?

STATEMENT OF THE CASE

I. THE '014 PATENT

U.S. patent 8,020,014, entitled, "Method for Power Reduction and a Device Having Power Reduction Capabilities," issued September 13, 2011 from an application filed November 9, 2007 claiming priority to a prior Patent Cooperation Treaty application filed May 11, 2005. Appx60. It generally discloses power reduction methods and devices that determine whether to power down a portion of a device, such as a cache memory, in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down. *Id.* (abstract).

In its "Background of the Invention," the '014 patent describes a goal of reducing power consumption in devices such as cache memories, particularly as they are used in mobile devices which are required to operate for long periods before being recharged. Appx64 (1:7-10, 2:13-21). It surveys prior techniques for achieving that goal, including, "selectively shutting down components," "reducing the clock frequency," and "dynamic voltage scaling (DVS)" involving "altering the voltage that is supplied to a processor … in response to the computational load demands." *Id.* (2:21-28). Among the prior techniques described and distinguished in the '014 patent background are two published U.S. applications, 2003/0145241 to Hu and 2002/0049918 to

Kaxiras, with Hu and Kaxiras being co-authors of the Hu article at issue in this case. *Id.* (2:61-67). Although not identical to Hu, these references describe the same process disclosed therein, namely, "powering down cache lines that were not accessed during a cache line decay interval." *Id*.

While these and other prior approaches share a general goal of power reduction, the approach they take differs markedly from that of the '014 patent. Most operate based on considerations of anticipated activity or performance. For example, the various "dynamic voltage scaling" techniques reduce voltage to system components "in response to the computational load demands." *Id.* (2:23-53). Similarly, the published applications of Hu and Kaxiras describe powering down portions of a cache "that were not accessed for a long time." *Id.* (2:65-67). By contrast, the '014 patent is the first disclosure of a technique for powering down portions of a cache based on whether doing so is estimated to actually save power. The evidence of record shows that, prior to the '014 patent, estimates of power gain and power loss associated with powering down portions of a cache were considered difficult to make. *See, e.g.*, Appx1791. ("[N]ot only are **direct comparisons of energy measurements hopeless** due to disagreement, but **even relative comparisons often fail to agree**. . . . [A]rchitectural description **simply doesn't contain an adequate amount of information to properly estimate power** . . .") (emphasis added). Hu itself

reinforces this difficulty in making power estimates. Appx888 ("Since both leakage and dynamic power values vary heavily with different designs and fabrication processes, it is difficult to nail down specific values for evaluation purposes."). Even Intel's own expert has acknowledged this difficulty, admitting that there are "multiple reasons" why leakage and dynamic power values could vary with specific design and fabrication processes. Appx1585-1586 (162:24-163:4).

The inventors of the '014 patent realized that so-called "dirty information," *i.e.*, data that appears in a portion of cache that has not yet been written back to main memory, presents a power cost if that portion of the cache is powered down. This is because, in order to not permanently lose that data, it must be written back to main memory before the portion of cache is powered down, and this writeback operation itself costs power. Appx65 (3:35-49), Appx66 (6:42-45). The inventors reasoned that, if the amount of dirty information is small, the power gain from powering down the portion of the cache outweighs the power loss from the few extra writeback operations. On the other hand, if the amount of dirty information is large, the power gain from powering down the portion of the cache may be less than the power loss from the large number of extra writeback operations. The '014 patent quantifies this crossover, where the estimated power gain overtakes the estimated power loss,

using a "power gating threshold TH," which is a threshold number of dirty bits that is compared to the actual number of dirty bits to make the determination of whether to power down. Appx65 (4:58-67).

II. THE APPLIED REFERENCES

A. Takahashi

U.S. patent 5,761,715 ("Takahashi") entitled, "Information processing device and cache memory with adjustable number of ways to reduce power consumption based on cache miss ratio," issued June 2, 1998 from an application filed August 8, 1996. Appx869. Though it suffers from an apparently poor translation from the original Japanese,[1] Takahashi does express a goal to reduce power consumption by powering down portions of a cache, and keeping them powered down so long as doing so does not adversely impact performance as indicated by an increase of the cache-miss rate. Takahashi repeatedly affirms its concern that reducing power not compromise the performance and speed of the device. *See*, *e.g.*, Appx875 (2:24-33) (seeking to avoid "decreasing the operation speed and the performance of the information processing device"); Appx876 (3:61-65) ("[T]he number of ways to be accessed is decreased … without decreasing the performance of the cache

[1] The Board itself acknowledges that "Takahashi's disclosure … unfortunately is not a model of clarity." Appx32.

memory."); Appx880 (12:12-17) ("[T]he present invention can reduce the power consumption without decreasing the performance …"). In particular, it seeks to avoid the impact on performance associated with an increase in the cache-miss rate. Appx878 (8:56-60) ("The pre-process operation must be executed in order to avoid the increasing of the cache-miss rate temporally when required data items are stored in the way that is eliminated …"); Appx879 (9:50-62) (describing how "the number of activated ways is changed from one way to two ways" when "the cache-miss rate is increased").

In furtherance of this goal, Takahashi's basic principle of operation is to preemptively cut power to portions of a cache, check the resulting change of the cache-miss rate, and use that change to determine whether to power those portions back up. Appx875 (2:19-21) ("[T]he number of ways to be accessed is dynamically switching while checking the change of the cache hit rate in memory access operations."). The operation flow in Figure 4 is informative:

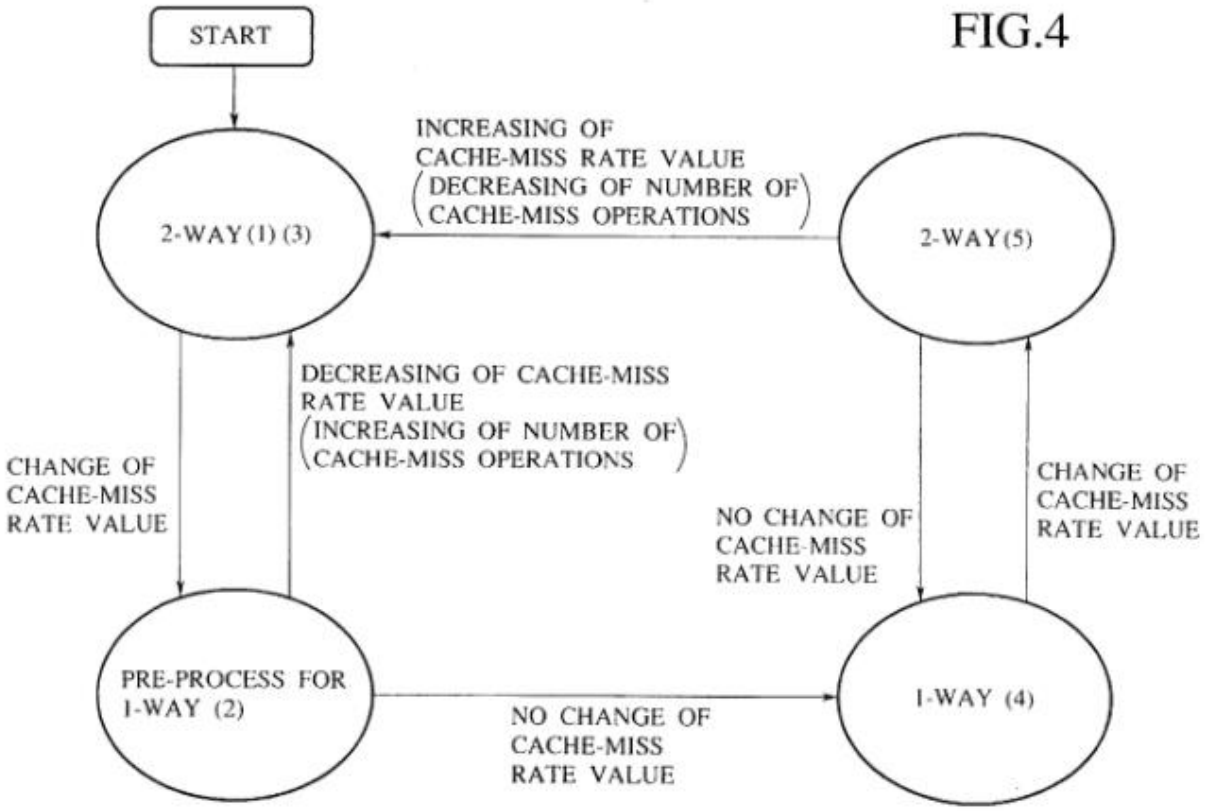

Appx873. Intel's allegations about Takahashi's disclosure focus on the evaluation of the change of cache-miss rate that occurs in the pre-process for 1-way mode at the bottom left. Appx105-111; Appx1940-1946. Takahashi begins its process in the 2-way mode at the top left, in which both cache ways are powered on, and automatically enters the pre-process for 1-way mode by powering down one of the cache ways. Appx878 (8:61-64) ("[T]he number of ways is decreased from two ways to one way. In this case, way 1 is eliminated from the access operation. . . . power source is disconnected from the way 1."). Only in this mode, when one of the two cache ways has already been powered down, does Takahashi compare the current cache-miss rate to the

"predetermined value" (a prior-measured cache-miss rate) to determine the change of cache-miss rate. *See* Appx876 (3:23-31) (describing "way number determination means" as "comparing the cache-miss rate at the present time … with the cache-miss rate at the previous time"). Takahashi states that, "if the cache-miss rate is increased [*i.e.* current value exceeds the predetermined value], the number of the activated ways is changed from one way to two ways," that is from "one way" in the pre-process mode (bottom left of Fig. 4) back to "two ways" (top left). Appx879 (9:59-61). On the other hand, "when the cache-miss rate does not exceed the predetermined value [*i.e.*, is not increased], the process flow of the cache memory 103 is shifted to the one way process," shown at the bottom of Fig. 4, in which the second cache way stays powered down. *Id.* (9:55-59). Thus, the comparison of the current cache-miss rate value with the predetermined value, *i.e.*, the evaluation of the change of cache-miss rate, occurs *after* one of Takahashi's two cache ways has already been powered down, as part of a determination of whether to power back up. In this way, Takahashi will check the impact of powering down portions of a cache on performance, and will keep those portions powered down so long as doing so does not detrimentally impact performance, including by increasing the cache-miss rate. Appx876 (3:61-65) ("[T]he number of ways to be accessed is decreased … without decreasing the performance of the cache memory.");

Appx878 (8:56-60) ("The pre-process operation must be executed in order to avoid the increasing of the cache-miss rate …").

B. Hu

"Let Caches Decay: Reducing Leakage Energy via Exploitation of Cache Generational Behavior," ("Hu") is a 2002 academic paper authored by Hu, Kaxiras, and Martonosi. As described in the background of the '014 patent, Hu describes performing simulations of "powering down cache lines that were not accessed during a cache line decay interval." *See* Appx64 (2:61-65); Appx886; Appx893. Hu describes simulating values of a quantity it calls the "normalized cache leakage energy," or just "normalized leakage energy." Appx888. Hu performs its simulations over "a wide[] range of decay intervals, from 1 K to 512 K cycles, to explore the design space thoroughly." Appx891.

In addition to the decay intervals, Hu also considers the effect of another independent variable on the normalized leakage energy: the *L2Access:leak* ratio. Appx897 ("We also wanted to explore the sensitivity of our results to different ratios of dynamic L2 energy versus static L1 leakage."). Hu explains that *L2Access* represents "energy per L2 cache access," where L2 is a level 2 cache, and *leak* is the "leakage energy dissipated by the L1 data cache" per clock cycle, where L1 is a level 1 cache. Appx889-890. In Figure 9, Hu plots the simulated values of normalized leakage energy as a function of decay

interval for four selected values of the *L2Access:leak* ratio parameter. Appx898 ("The curves correspond to *L2Access:leak* ratios of 5, 10, 20, and 100.").

III. THE PROCEEDINGS BELOW

A. Intel's Grounds of Alleged Invalidity

Each of Intel's two *inter partes* review petitions addressed one of the two independent claims of the '014 patent: claim 12 in IPR2018-01312, and claim 1 in IPR2018-01661. Each alleged two grounds of invalidity against the challenged independent claim: (1) obviousness over Takahashi, and (2) obviousness over Takahashi and Hu. Appx101-111; Appx118-129; Appx1936-1946; Appx1948-1958. Intel made essentially the same arguments about these references in relation to the challenged claims across the two petitions.

For ground (1), Intel pointed to the transition from Takahashi's "pre-process for one way process" to its "one way process" as allegedly teaching the claimed "selectively providing power to at least a portion of a component," arguing that in this transition, "the power supply to a specific cache way … is selectively turned off." Appx102-105; Appx1937-1940. As discussed above, Takahashi's pre-process for one way process already has the power supply to a specific cache way turned off, and Intel's position appears to have been based on a mistaken belief that both cache ways are still powered on in this pre-process mode, a belief that Intel maintained throughout the proceeding below.

See Appx325; Appx2196 ("Takahashi expressly discloses that the entire cache is powered on during the 'pre-process' mode …"). For the claimed, "determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down," Intel pointed to Takahashi's comparison of the current value of the cache-miss rate with a "predetermined value." Appx105-111; Appx1940-1946. Relying on the '014 patent disclosure for guidance, Intel argued by analogy that Takahashi's predetermined value must serve the same purpose as the "power gating threshold TH" of the '014 patent, and thus be representative of an estimated power gain resulting from powering down a portion of the cache. Appx107-111; Appx1942-1946. As discussed above, Takahashi's "predetermined value" is actually just a prior measured value of the cache-miss rate, which is compared to the current value to determine the *change* of the cache-miss rate. *See* Appx875-876 (2:38:46, 3:23-31) (describing "the way number determination means" as comprising a "first comparison means for comparing the cache-miss rate at the present time … with the cache-miss rate at the previous time").[2]

[2] The correct understanding of Takahashi's "predetermined value" was affirmed by this Court as well. *Intel Corp. v. VLSI Tech. LLC*, 858 F. App'x 349, 352 (Fed. Cir. 2021) ("Takahashi teaches storing a previous measurement of the cache-miss rate as a 'predetermined value.'").

For ground (2), Intel maintained that Takahashi taught all elements of the claim, cross-referencing its arguments from ground (1). However, it also added arguments that Hu taught the claimed, "determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down." *See* Appx1024, Appx1953 ("Hu uses the *L2Access:leak* ratio to determine the relationship between estimated power loss (numerator) and estimated power gain (denominator), and determines when to power down a cache line based on that relationship."). Intel's argument focused on simulations of the normalized leakage energy as a function of decay interval and *L2Access:leak* ratio, in particular those illustrated in Hu's Figure 9. Appx4050-4052, Appx5042-5044; *see also* Appx898 (Fig. 9). Intel pointed to these to argue that Hu taught how to determine a decay interval based on the *L2Access:leak* ratio, the latter of which it contended corresponds to the claimed "relationship between an estimated power gain and an estimated power loss resulting from powering down." *Id*.

B. The Board's Original Final Written Decision

With regard to ground (1), obviousness over Takahashi, the Board noted in its 2020 Original Final Written Decision,

> Patent Owner raises two main arguments . . . . that Takahashi's comparison between a cache-miss rate and a predetermined value is not a "relationship between an estimated power gain and an estimated power loss" . . . . [and] that Takahashi does not "determine whether to power down" . . . because when Takahashi's device compares a cache-miss rate to a predetermined value, it already has powered down one cache way.

Appx571, Appx2504. The Board declined to reach the second issue "[b]ecause the first issue is dispositive." *Id.* There, the Board found that "Patent Owner's cited evidence amply supports" a factual finding that Takahashi's predetermined value is nothing more than a previously measured cache-miss rate, which is used in comparison with a current measured value to determine the change of the cache-miss rate. Appx573, Appx2505.

With regard to ground (2), the Board found that "Hu discloses determining values of the normalized cache leakage energy for different decay intervals at four values of the *L2Access:leak* ratio—5, 10, 20, and 100." Appx580, Appx2512. Dismissing Intel's allegation "that Hu teaches that the higher the *L2Access:leak* ratio, the longer the system should wait before

powering down," the Board recognized that, in fact, "Hu does not describe using the *L2Access:leak* ratio to determine an appropriate decay interval, nor … to determine whether to power down a portion of a cache." Appx580, Appx2513. As noted below, this Court reversed the Board's finding on this issue.

The Board further acknowledged that Patent Owner raised "several arguments in response to Petitioner's contentions," including "that Hu does not teach a relationship between an estimated power gain and … power loss because Hu's simulated system focuses on ratios of values," and also that, "the *L2Access:leak* ratio is a static, intrinsic device characteristic that … does not reflect estimated power gain and loss resulting from powering down." Appx579, Appx2511-2512. However, it declined to rule on these, choosing instead to "focus on one argument we find particularly persuasive—that Hu does not teach using the *L2Access:leak* ratio to 'determine whether to power down' a portion of a cache." *Id.*

Additionally, the Board found that the *L2Access:leak* ratio constituted "assumed values . . . as inputs to generate simulation," which VLSI had argued rendered the ratio distinct from the claimed "estimated power gain and an estimated power loss resulting from powering down." Appx580, Appx2512.

C. The Federal Circuit's Decision

Intel appealed the Board's original Final Written Decisions to this Court, arguing (1) obviousness over Takahashi alone based on its understanding that Takahashi's "predetermined value was 'representative of the estimated power gain' from powering down a cache way"; (2) that Intel's disclosure of its theory combining Takahashi with Hu for particular claim limitations was not untimely; and (3) that the Board erred in finding that Hu did not teach determining whether to power down portions of a cache based on its *L2Access*:*leak* ratio. *Intel Corp. v. VLSI Tech. LLC*, 858 F. App'x 349, 352-355 (Fed. Cir. 2021). This Court affirmed the Board's determination that Takahashi does not disclose "an estimated power gain," but reversed the Board's decision that Intel's disclosure of its obviousness theory was untimely as well as its determination that Hu did not teach using the *L2Access:leak* ratio to determine whether to power down a portion of a cache. *Id.* Accordingly, this Court remanded to the Board "to address whether the combination of Takahashi and Hu satisfies all of the claim limitations, and whether there was a motivation to combine Takahashi and Hu with a reasonable expectation of success." *Id.*

D. The Board's Final Written Decision on Remand

The Board acknowledged this Court's instruction to determine whether the combination of Takahashi and Hu satisfies all of the claim limitations, and

whether there was a motivation to combine Takahashi and Hu with a reasonable expectation of success, in view of this Court's determination that Hu taught determining whether to power down portions of a cache based on the *L2Access:leak* ratio. Appx4-5. With regard to the claimed, "determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down," the Board found that, "Hu describes a 'relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component' as claimed." Appx27. In support, the Board argued that "Hu explicitly describes using estimates for the L1 static leakage energy and L2 access energy," citing values of "3 to 5 nJ" for the former and "0.45 nJ" for the latter. *Id.* The Board further found a motivation to combine the references, discounting Patent Owner's argument that Takahashi's pre-process for 1-way mode taught away from Hu's decay interval approach. Appx30-34. The Board also found that there would be a reasonable expectation of success in making the combination to reach the challenged claims, noting that "both references teach using information related to cache misses to determine when to power down a portion of a cache." Appx36.

SUMMARY OF ARGUMENT

The Board made four errors, each of which is an independent ground for vacatur and remand.

First, the Board erred by relying on quantities in Hu, called *L2Access* and *leak*, which do not and cannot correspond to any amount of power gain or loss "*resulting from powering down the at least portion of the componen*t," as recited by the challenged independent claims. Both represent static, intrinsic quantities (*per* access and *per* clock cycle, respectively), as explicitly stated in Hu and acknowledged by the Board, which carry no information about the number of accesses or number of clock cycles. The missing information is essential to determine power gain or loss "resulting from powering down," and thus the Board's finding that Hu teaches these limitations in the absence of that essential information is clear error.

Second, the Board erred by relying on calculated numerical values for *L2Access* ("3 to 5 nJ per access") and *leak* ("0.45 nJ" per cycle) for the claimed "**estimated power gain**" and "**estimated power loss**," but then relying on entirely different, improvised values (5, 10, 20 and 100) for the ratio of these quantities employed in connection with Hu's simulation of normalized leakage energy for the claimed "determining whether to power down … in response to a relationship between **an estimated power gain** and **an estimated power**

loss." By dissecting the claim and relying on disconnected quantities from the reference to reach the claim, the Board exercised improper hindsight in reaching the conclusion of obviousness.

Third, the Board erred in finding a motivation to combine Takahashi with Hu by misapprehending Takahashi's basic principle of operation as set forth in the reference (even holding acknowledged ambiguities of the reference's teaching against the validity of the challenged patent), resulting in a failure to recognize clear teaching away from the proposed combination.

Fourth, the Board erred by failing to make an analysis, supported by substantial evidence, of the requisite reasonable expectation of success to support the ultimate conclusion of obviousness. At most, for a subset of claims, the Board provided only a cursory and conclusory assertion based on an alleged similarity between highly generalized characterizations of the references.

STANDARD OF REVIEW

The Court reviews final written decisions of the Board in accordance with the Administrative Procedures Act, and sets aside those decisions if they are "'arbitrary, capricious, an abuse of discretion, or otherwise not in accordance with law.'" *HTC Corp. v. Cellular Communications Equipment, LLC*, 877 F.3d 1361, 1367 (Fed. Cir. 2017); *Vicor Corp. v. SynQor, Inc.*, 869

F.3d 1309, 1320 (Fed. Cir. 2017) (quoting 5 U.S.C. § 706(2)). Legal conclusions of obviousness are reviewed *de novo*, with underlying factual findings reviewed for substantial evidence. *HTC Corp.*, 877 F.3d at 1369. Obviousness is a question of law with underlying factual findings relating to the scope and content of the prior art; differences between the prior art and the claims at issue; the level of ordinary skill in the pertinent art; the presence or absence of a motivation to combine or modify prior art with a reasonable expectation of success; and any objective indicia of non-obviousness. *Acoustic Tech., Inc. v. Itron Networked Sols., Inc.*, 949 F.3d 1366, 1373 (Fed. Cir. 2020). "Substantial evidence review asks 'whether a reasonable fact finder could have arrived at the agency's decision and requires examination of the record as a whole, taking into account evidence that both justifies and detracts from an agency's decision.'" *Intelligent Bio Systems, Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1366 (Fed. Cir. 2016) (quoting *In re Gartside*, 203 F.3d 1305, 1312 (Fed. Cir. 2000)). Substantial evidence is such relevant evidence as a reasonable mind might accept as adequate to support a conclusion. *B/E Aerospace, Inc. v. C&D Zodiac, Inc.*, 962 F.3d 1373, 1379 (Fed. Cir. 2020) (citing *Consolidated Edison Co. v. NLRB*, 305 U.S. 197, 229 (1938))

ARGUMENT

I. THE BOARD ERRED IN RELYING ON HU'S *L2ACCESS* AND *LEAK* PARAMETERS FOR THE CLAIMED "ESTIMATED POWER GAIN AND ESTIMATED POWER LOSS RESULTING FROM POWERING DOWN THE AT LEAST PORTION OF THE COMPONENT"

The Board's conclusion of obviousness rests on its identification of *L2Access* and *leak* as the claimed "an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component." Appx27. But *L2Access* and *leak*, being intrinsic quantities that contain no information about operational conditions, are logically incapable of qualifying as estimated quantities "resulting from powering down." Instead, further information about the number of L2 accesses occurring upon shutdown and the amount of time of the shutdown (neither of which is provided or estimated in Hu), is necessary to convert these quantities into energies "resulting from powering down." Accordingly, the Board's finding of obviousness on this basis is not supported by substantial evidence.

The challenged independent claims of the '014 patent recite a determination of "whether to power down" a portion of a component " in response to a relationship between an **estimated power gain and an estimated power loss resulting from powering down** the at least portion of the component." Appx67 (7:31-35; 8:13-18). In accordance with this Court's prior ruling, the Board acknowledged "that Hu teaches using the *L2Access:leak* ratio

to determine whether to power down a portion of a cache." Appx20. The Board also identifies *L2Access* and *leak*, respectively, as the claimed "estimated power gain and estimated power loss resulting from powering down." Appx27 ("[T]he denominator of the *L2Access:leak* ratio reflects the static leakage power that would be saved by turning off a portion of the L1 cache. . . . [T]he numerator of the *L2Access:leak* ratio reflects the estimated amount of dynamic energy dissipated per L2 cache access … that would result from turning off a portion of the L1 cache.").

It is undisputed that *L2Access* and *leak* are intrinsic quantities that characterize the properties of Hu's device. *See, e.g.*, Appx889 ("The *L2Access:leak* ratio relates dynamic energy due to **an [*i.e.*, one] additional miss** (or writeback) to a **single clock cycle** of static leakage energy in the L1 cache.") (emphasis added); Appx889-890 (characterizing *L2Access* as "energy per cache access" and *leak* as "energy per cycle"); Appx16 (characterizing *leak* as "energy dissipated by the L1 data cache during one clock cycle") (internal quotations omitted); Appx23 (characterizing *L2Access* as "dynamic energy per L2 cache access").

L2Access and *leak* are thus not extrinsic quantities that characterize the "power … resulting from powering down," which would require additional information about the operation of the device upon shutdown, specifically the

number of L2 accesses that would occur upon powering down (for *L2Access*), and the amount of time (*i.e.*, number of clock cycles) for which the L1 cache would be powered down (for *leak*). Appx375, Appx478, Appx2246. This would be analogous to a claim reciting an "estimated weight" and a prior art reference solely disclosing a density. Density is an intrinsic quantity (characteristic of a material) and, while it may relate to weight (an extrinsic quantity) together with the volume of material, there is simply no way to actually obtain an "estimated weight" based on density alone. In fact, the weight could take literally any value depending on the volume of material. Similarly here, the power loss "resulting from powering down" could take literally any value for a given value of *L2Access*: if powering down required no writebacks (*e.g.*, if there were no dirty bits in the L1 memory), the power lost would be zero independent of the *L2Access* value, and if many writebacks were required, the power loss could be arbitrarily large for the same *L2Access* value. Similarly, the power gain resulting from powering down could take any value for a given value of *leak*: if powering down occurred for a very brief period of time, the power gain would be arbitrarily small, while if powering down occurred for a very long period of time, the power gain would be arbitrarily large. Simply put, neither *L2Access* nor *leak* contains enough information to

estimate power gain or power loss "resulting from powering down" a portion of a device, as the claims require.

The Board makes no acknowledgment or explanation for this critical shortcoming of the *leak* parameter, never pointing to any evidence or providing any analysis to suggest that Hu somehow takes into account the amount of time that L1 cache would be powered down in order to be able to calculate a power gain "resulting from powering down" that portion of the cache. Instead, it states in conclusory fashion, "Hu generally teaches … an estimated power gain (*e.g.*, static L1 cache leakage power saved by turning off a cache line) …" Appx41 (citing Appx888-891, Appx898-898). The portions of Hu cited here do not support the assertion, instead merely discussing Hu's "0.45 nJ static leakage **per cycle** estimate" for *leak* based on its literature review, and the results of its simulation as a function of decay interval for different values of the *L2Access:leak* ratio (Fig. 9). Being devoid of any discussion of a time period for which the L1 cache would be shut down, Hu (and the Board's analysis of it) omits a necessary parameter required to allow the intrinsic *leak* quantity to be converted into a power loss resulting from powering down.

As for *L2Access*, the Board actually acknowledges that it is an intrinsic quantity that does not carry the extrinsic information about the number of L2 accesses necessary to determine a power gain resulting from powering down.

Appx41 ("We agree with Patent Owner that the *L2Access* parameter itself does not reflect the number of L2 accesses that might be caused by turning off a cache line because it represents the dynamic power *per access* to L2 memory") (original emphasis). However, it proceeds to allege that, "Hu also generally teaches … an estimated power loss … [that] takes into consideration the additional L2 accesses that powering down a cache line containing dirty information would cause." Appx41-42 (citing Appx888-889, Appx1964-19965, Appx854). None of the evidence cited by the Board here supports its assertion. The cited portions of the Petition and accompanying expert declaration (which recite identical language) merely note the definition of *L2Access* from Hu as "dynamic energy due to an additional miss (or writeback)"; in other words, the same intrinsic energy *per access* to L2 memory. Appx1964-1965 (citing Appx889), Appx854. And the portions of Hu itself which are cited by the Board also do not support the notion that the reference teaches determining whether to power down based on an estimated power loss resulting from powering down. At most, the cited evidence indicates that the formula for normalized leakage energy can include a product of *L2Access* with a quantity called *extraL2Accesses*. Appx889. But this is also clearly inadequate because (1) there is no estimate of *extraL2Accesses*

anywhere in Hu,[3] and thus no estimated power loss, and (2) *extraL2Accesses* has not been found by either the Board or this Court to form the basis of any determination of whether to power down. *Intel Corp. v. VLSI Tech. LLC*, 858 F. App'x 349, 355 (Fed. Cir. 2021) ("Hu discloses using the *L2Access:leak* ratio to determine whether to power down a cache (or portion thereof) …"); Appx20 ("[W]e consider the parties' arguments based on the complete record and in view of (1) the Federal Circuit's determination that Hu teaches using the *L2Access:leak* ratio to determine whether to power down a portion of a cache …").

As neither *L2Access* nor *leak* can possibly constitute "an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component," as recited by the challenged independent claims, the Board's contrary findings lack substantial evidence support and should be reversed. *PAR Pharm., Inc. v. TWI Pharms., Inc.*, 773 F.3d 1186, 1194 (Fed. Cir. 2014) (obviousness requires the challenger to "prove that all claimed limitations are disclosed in the prior art"); *Wireless Protocol*

[3] Moreover, the Board has relied solely on computed values of *L2Access* and *leak* themselves for the claimed "estimated power gain" and "estimated power loss." Appx27 ("Hu explicitly describes using estimates for the L1 static leakage energy and L2 access energy … dynamic energy per L2 access in the range of 3 to 5 nJ … 0.45 nJ static leakage per cycle estimate obtained from literature review.") (cleaned up).

Innovations, Inc. v. TCT Mobile, Inc., 771 F. App'x 1012, 1013 (Fed. Cir. 2019) (reversing finding of obviousness where "the combination of asserted prior-art references does not disclose every element of the challenged claims").

II. THE BOARD ALSO ERRED IN RELYING ON HU'S COMPUTED VALUES OF *L2ACCESS* AND *LEAK* THAT WERE NOT USED IN ITS SIMULATION OF NORMALIZED LEAKAGE ENERGY

The Board further erred in adopting one set of estimated values for *L2Access* and *leak* from Hu for the claimed, "**an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component**," and an unrelated set of non-estimated values for the *L2Access:leak* ratio for the same limitation in connection with the claimed, "determine whether to power down at least the portion of the component during a low power mode in response to a relationship between **an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component**." This inconsistent approach ignores a basic inadequacy of Hu: no estimated value of *L2Access*, *leak*, or their ratio *L2Access:leak* is ever used in connection with any determination of whether to power down portions of a device, as required by challenged independent claims 1 and 12. *See University of Strathclyde v. Clear-Vu Lighting, LLC*, 17 F.4th 155, 160-62 (Fed. Cir. 2021) (finding Board's conclusion of obviousness

unsupported by substantial evidence where claim element "neither taught nor suggested by the prior art of record.").

For the claimed "estimated power gain" and "estimated power loss," the Board explicitly relies upon specific numerical values quoted in Hu that are derived from a literature review. Appx17 ("Hu estimates *L2Access* to be in the range of 3 to 5 nJ per L2 access … and *leak* to be in the range of 0.45 nJ per clock cycle …") (citing Appx889-890). From these estimates, Hu arrives at a general estimate of 8.9 for the *L2Access:leak* ratio in a typical system of interest. Appx890 ("[W]e get a ratio of 8.9 relating extra miss power to static leakage per cycle."); Appx17 ("Hu estimates an *L2Access:leak* ratio of 8.9 for devices using caches of those sizes."). The determination of "whether to power down at least the portion of the component during a low power mode in response to a relationship between an estimated power gain and an estimated power loss" should presumably then depend upon these values, or otherwise the Board should provide some reasoned explanation supported by evidence to justify using some other values. But neither of these is the case.

Instead, the Board simply pivots to a different portion of Hu that describes "determining the normalized cache leakage energy across a set of benchmark applications for different values of the *L2Access:leak* ratio, including ratios of 5, 10, 20, and 100." Appx17 (citing Appx888-889,

Appx895-898, Figs. 8, 9). This Court also pointed to the same disclosure of Hu, notably Fig. 9, as teaching a determination of whether to power down based on the *L2Access:leak* ratio. *Intel Corp. v. VLSI Tech. LLC*, 858 F. App'x 349, 355 (Fed. Cir. 2021) ("Hu's experimental results also show that the *L2Access:leak* ratio affects the optimal decay interval for powering down a cache. . . That is, Hu teaches using the *L2Access:leak* ratio to estimate the optimal decay interval, which in turn is used to determine when to power down a cache.") (citing Appx898). As shown below, Hu's Fig. 9 plots the normalized leakage energy as a function of decay interval for exactly four values of *L2Access:leak*: 5, 10, 20, and 100.

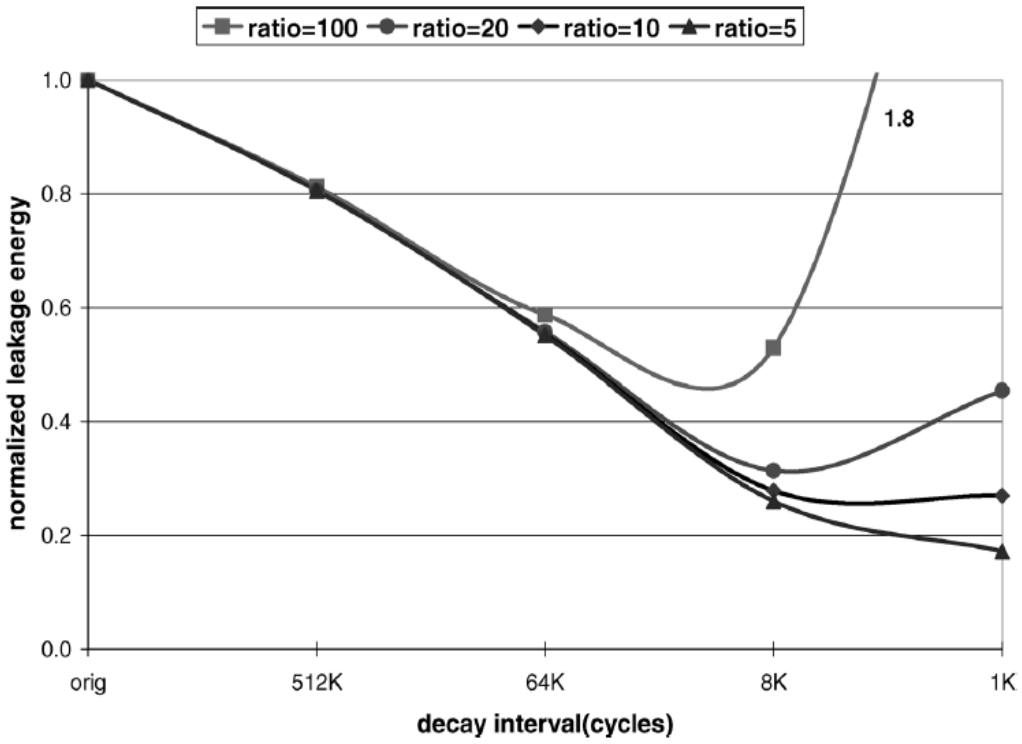

Fig. 9. Normalized L1 data cache leakage energy averaged across SPEC suite for various *L2Access:leak* ratios.

Appx898. Although this Court held that Fig. 9 shows "that the *L2Access:leak* ratio affects the optimal decay interval for powering down a cache," there is no connection to the "estimate[d] … *L2Access:leak* ratio of 8.9," let alone the estimates of "*L2Access* to be in the range of 3 to 5 nJ per L2 access … and *leak* to be in the range of 0.45 nJ per clock cycle." *Intel Corp*, 858 F. App'x at 355; Appx17. Instead, it employs selected, independent values of 5, 10, 20 and 100. *See also* Appx2510 (prior Final Written Decision of the Board noting Hu "determines the normalized cache leakage energy … for different **assumed** values of the *L2Access:leak* ratio … 5, 10, 20, and 100.") (emphasis added).

In effect, the Board has dissected the claim, pointing to one aspect of Hu's disclosure for the "estimated power gain" and "estimated power loss," and a wholly different aspect of Hu's disclosure bearing no discernible relationship to the first for the claimed determination of "whether to power down at least the portion of the component … in response to a relationship between an estimated power gain and an estimated power loss." Such dissection of the claim and piecemeal reconstruction of the prior art to arrive at the claimed invention is improper and reflects hindsight bias. *See, e.g., Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 36 (1966) (cautioning against "the temptation to read into the prior art the teachings of the invention in issue"); *In re Gulack*, 703 F.2d 1381, 1385 (Fed. Cir. 1983) ("Under section

103, the board cannot dissect a claim … The claim must be read as a whole."); *Application of Kamm*, 452 F.2d 1052, 1056-57 (CCPA 1972) (noting the existence of "a basis mandate inherent in § 103–that a piecemeal reconstruction of the prior art patents in the light of appellants' disclosure shall not be the basis for a holding of obviousness.") (cleaned up); *Application of Wesslau*, 353 F.2d 238, 241 (CCPA 1965) ("It is impermissible within the framework of section 103 to pick and choose from any one reference only so much of it as will support a given position …").

As the Board's findings for the claimed "estimated power gain" and "estimated power loss" rest on teachings of Hu that are disconnected from the teachings upon which its findings are based for the claimed determination of "whether to power down at least the portion of the component during a low power mode in response to a relationship between an estimated power gain and an estimated power loss," the Board's conclusion of obviousness is not founded upon substantial evidence and should be reversed.

III. THE BOARD'S FINDING OF A MOTIVATION TO COMBINE IS NOT SUPPORTED BY SUBSTANTIAL EVIDENCE

Among this Court's imperatives on remand was for the Board to determine "whether there was a motivation to combine Takahashi and Hu with a reasonable expectation of success." *Intel Corp.*, 858 F. App'x at 355. The Board's finding of motivation to combine Takahashi and Hu, as applied to all

challenged claims of the ’014 patent, misapprehends a critical incompatibility between the references. While Hu determines whether to power down a portion of a cache based on an elapsed period of time (which period is, according to this Court’s finding, determined by the *L2Access:leak* ratio), Takahashi never makes a determination of whether to power down a portion at all. Instead, its basic principle of operation is to preemptively power down a portion and check the resulting change of cache-miss rate to determine whether to stay powered down (if the rate has not increased) or to power back up (if the rate has increased). The Board errs in two important respects. First, while acknowledging ambiguities in Takahashi, it holds the ambiguity against the validity of the patent while simultaneously misapprehending clear disclosures in Takahashi demonstrating its principle of operation. Second, it discounts Takahashi’s clear goal of avoiding increasing the cache-miss rate, and argues that Takahashi’s “pre-process operation … achieves Takahashi’s stated purpose,” such that “modifying Takahashi to incorporate Hu’s teaching” would not frustrate Takahashi’s “stated purpose.” The Board’s conclusion does not logically follow from its premise, and the Board’s over-generalized definition of Takahashi’s principle of operation is unjustified and fails to overcome the clear fact that Hu’s “decay interval” technique would upend Takahashi’s fundamental principle of operation, rendering the combination non-obvious.

Plas-Pak Indus., Inc. v. Sulzer Mixpac AG, 600 F.App'x 755, 760 (Fed. Cir. 2015); *see also Tokai Corp. v. Easton Enters., Inc.*, 632 F.3d 1358, 1378 (Fed. Cir. 2011).

Takahashi's fundamental principle of operation is illustrated in its Fig. 4:

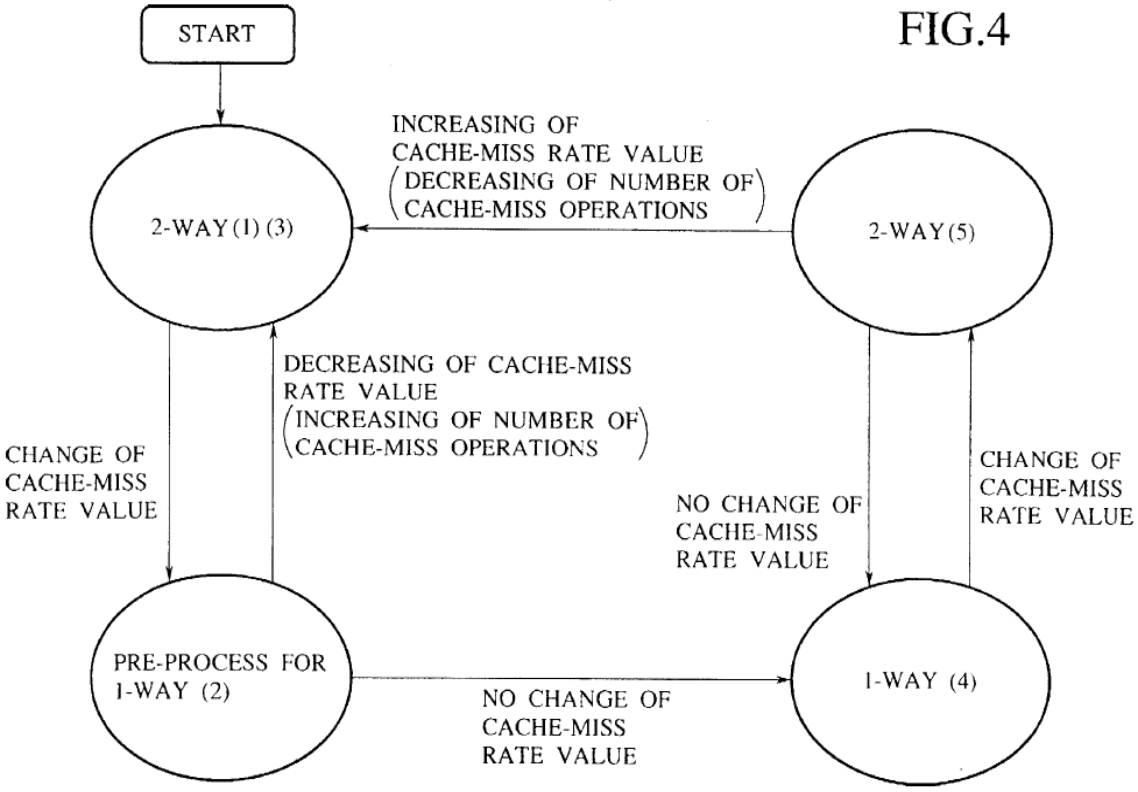

Appx873. Takahashi begins its operation, at the top left oval, in the "2 way operation mode" at (1), in which two cache ways are powered. Appx879 (9:41-43). The operation then enters the "pre-process for 1-way" mode at (2), in which "the number of ways is decreased from two ways to one way" because "the power source is disconnected from the way 1." Appx 878-879 (8:61-64, 9:43-49); *see also* Appx1560 (137:19-25) ("Q. Do you agree that in

Takahashi's preprocess operation the power source is first disconnected from one of the cache ways? A. (Witness reviews document) Yes. The patent states the power source is disconnected from way one."). The transition from two ways being powered in the 2-way operation mode to only one way being powered in the pre-process for 1-way operation mode occurs automatically based on the passage of a fixed, "predetermined time period" ("1024 memory access operations"). Appx879 (9:43-49); Appx1561 (138:11-21) ("Q. Okay. How does Takahashi's system determine whether to shift from two-way operation to the preprocess for one-way option – operation? A. As stated on line, I believe it's 43 here … in column 9 – 'a predetermined time period is elapsed (here the time period means … 1024 memory access operations) the operation is shifted to the preprocess for one-way operation.'"); Appx1387-1389 (noting shutdown of cache way "occurs automatically based solely on the passage of a fixed, 'predetermined time period'"). Thus, where Takahashi describes "the cache-miss rate does not exceed the predetermined value," it is describing a transition between two identically-powered modes (pre-process for 1-way, and 1-way) that both have only the same single cache way powered. Appx879 (9:50-62).

It makes sense that Takahashi's system powers down cache ways first, and only then evaluates the impact on cache miss rate, because its basic

principle of operation requires the cache way to be powered down first in order to produce a measurable effect on cache-miss rate via the "cache-miss rate measuring circuit 140" (with the goal "to avoid the increasing of the cache-miss rate"). Appx871, Appx878-879 (Fig. 2, 8:56-60, 9:50-62); see also Appx1560 (137:8-9) ("I would say that a goal of Takahashi is to avoid increasing the cache miss rate.").

The Board acknowledges that Takahashi is "unfortunately … not a model of clarity." Appx32. In particular, on the issue of whether it shuts down portions of a cache before or after checking the change of cache-miss rate, the Board "find[s] it ambiguous on this point, *i.e.*, whether it fully powers down a cache way before measuring a change in the cache-miss rate." Appx33. Nevertheless, though it remains Petitioner's burden to prove obviousness, the Board proceeds to hold the ambiguity in Takahashi *against* the validity of the challenged patent. Appx33 ("Because we find that Takahashi does not unambiguously teach powering down before determining a change in the cache-miss rate, we do not agree with Patent Owner's argument."); *cf. Dynamic Drinkware, LLC v. National Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015) ("In an inter partes review, the burden of persuasion is on the petitioner to prove 'unpatentability by a preponderance of the evidence,' 35 U.S.C. § 316(e), and that burden never shifts to the patentee."); *In re Magnum*

Oil Tools Int'l, Ltd., 829 F.3d 1364, 1376 (Fed. Cir. 2016) (holding that "no burden shifts from the patent challenger to the patentee" for issues of "what the prior art discloses, whether there would have been a motivation to combine the prior art, and whether that combination would render the patented claim obvious."). Moreover, though Takahashi is ambiguous in many ways, it clearly teaches a principle of operation that *requires* preemptive shut down of a cache way to induce a change of cache-miss rate that is used to determine whether to stay powered down or power back up. Appx876 (3:61-65); Appx879 (9:51-61).

The Board further erred in misapprehending the goal of Takahashi, stating in a footnote of its Final Written Decision that, "[w] also do not agree with Patent Owner's unsupported assertion in its Sur-reply that Takahashi's stated purpose is to avoid increasing the cache-miss rate." Appx33 (n. 14). In fact, as shown above and demonstrated in Patent Owner's IPR briefing, Takahashi is explicit (and even Petitioner's expert agrees) that its focus is "to avoid the increasing of the cache-miss rate." *See* Appx2141 (discussing Takahashi's goal to avoid increasing cache-miss rate, with citations to Takahashi and Petitioner's expert testimony); Appx2236 (describing Takahashi's principle of operation "to 'avoid the increasing of the cache-miss rate temporally when required data items are stored in the way that is eliminated.'") (citing Appx878-879 (8:56-60, 9:50-55)).

To replace Takahashi's approach, of preemptively shutting down a cache way and checking for a change of cache-miss rate to determine whether to power back up, with Hu's approach, of determining when to shut down based on a decay interval determined according to the *L2Access:leak* ratio, would eviscerate Takahashi's essential aspect and leave behind nothing of its principle of operation, rendering it incapable of fulfilling its goal to avoid increasing the cache-miss rate. *Univ. of Maryland Biotech. Inst. v. Presens Prevision Sensing GmbH*, 711 F. App'x 1007, 1010 (Fed. Cir. 2017) ("[A] person of ordinary skill generally would not be motivated to modify a reference by contradicting its basic teachings … or by making it 'inoperable for its intended purpose.'"). As such, the combination of Takahashi and Hu would not be obvious, and the Board's finding to the contrary is not supported by substantial evidence and should be reversed on this basis alone. *Plas-Pak Indus.*, 600 F. App'x at 759-60; *In re Gordon*, 733 F.2d 900, 902 (Fed. Cir. 1984) (finding that a modification which renders the invention inoperable for its intended purpose is not obvious because it teaches away from the invention); *Application of Ratti*, 270 F.2d 810, 813 (CCPA 1959) (finding nonobviousness where "suggested combination of references would require a substantial reconstruction and redesign of the elements" of a reference and "change … the basic principle under which … [it] was designed to operate").

The Board attempts to backstop its conclusion by arguing that "regardless of whether Takahashi powers down a cache way before or after measuring the cache-miss rate," the combination with Hu would still not render it unsuitable for its stated purpose. Appx33. The basis of the Board's argument is that "Takahashi's pre-process operation determines the optimal number of cache ways based on changes in the cache-miss rate resulting from powering down," and "[t]his determination achieves Takahashi's stated purpose of reducing power consumption without negatively impacting the information processing device's performance." *Id.* But this argument must also fail. The "pre-process operation" which "achieves Takahashi's stated purpose" is precisely the portion of Takahashi that would be replaced by Hu's approach, namely the preemptive shutting down of a cache way to check the change of cache-miss rate. Appx743-748, Appx838-844. In other words, in arguing that Takahashi's stated purpose would remain fulfilled in combination with Hu, the Board points to the very portion of Takahashi that would be eliminated in the combination. Thus, again here, the Board's conclusion is unsupported by substantial evidence and should be reversed.

IV. THE BOARD'S FINDING OF A REASONABLE EXPECTATION OF SUCCESS IS NOT SUPPORTED BY SUBSTANTIAL EVIDENCE

It is well-established under Federal Circuit law that, to invalidate a patent claim based on obviousness, a challenger must demonstrate "[1] that a

skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and [2] that the skilled artisan would have had a reasonable expectation of success in doing so." *ActiveVideo Networks, Inc. v. Verizon Commc'ns, Inc.*, 694 F.3d 1312, 1327 (Fed. Cir. 2012) (cleaned up); *see also Eli Lilly and Co. v. Teva Pharms. Int'l GmbH*, 8 F.4th 1331, 1344 (Fed. Cir. 2021) ("[W]e must first emphasize the clear distinction in our case law between a patent challenger's burden to prove that a skilled artisan would have been motivated to combine prior art references and the additional requirement that the patent challenger also prove that the skilled artisan would have had a reasonable expectation of successfully achieving the claimed invention from the combination."); *Intelligent Bio-Systems, Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367 (Fed. Cir. 2016) (noting "motivation to combine" and "reasonable expectation of success" are "two different legal concepts"). The Board's decision fails to provide any significant analysis, supported by substantial evidence, demonstrating that a skilled artisan would have had a reasonable expectation of success in achieving the claimed invention based on the combinations of references upon which it relies. Instead, in most instances, the Board makes nothing more than conclusory assertions unsupported by any kind of reasoned analysis or evidence. In the two instances in which it provides anything further, it either relies on similarly

conclusory assertions by reference to Petitioner's IPR briefing, or on its own grossly overgeneralized characterization of the references to allege their similarity as a basis supporting reasonable expectation of success. In all cases, the Board's finding lacks support by substantial evidence.

The Board acknowledges that "[a]n obviousness determination requires finding 'both that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.'" Appx10 (citing *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367-68 (Fed. Cir. 2016)). Yet, for all findings of obviousness for dependent claims of the '014 patent over the combination of Takahashi and Hu, it provides no analysis at all and cites no evidence supporting a finding that a skilled artisan would have had a reasonable expectation of success in reaching the claimed invention. Instead, the Board simply states, in conclusory fashion, that "a person of ordinary skill in the art would have combined the teachings of Takahashi and Hu in the manner asserted with a reasonable expectation of success." Appx38 (claims 2, 3, 13, and 14); Appx42 (claims 4, 5, 15, and 16); Appx45 (claims 18 and 20).[4] This is plainly insufficient as a

[4] Notably, the Petitioner's IPR briefs also appear to be devoid of any showing, or even mention, of a reasonable expectation of success for the specific proposed combination of Takahashi and Hu. Appx74-147, Appx323-

matter of law. *In re Kahn*, 441 F.3d 977, 989 (Fed. Cir. 2006) ("[R]ejections on obviousness grounds cannot be sustained by mere conclusory statements; instead, there must be some articulated reasoning with some rational underpinning to support the legal conclusion of obviousness."). Nor can the Board's analysis for the independent claims, which itself is cursory and inadequate as detailed below, support the same finding for the dependent claims which recite additional limitations. *See Gillette Co. v. SC Johnson & Son, Inc.*, 919 F. 2d 720, 724 (Fed. Cir. 1990) (obviousness "requires analysis of a claimed invention *as a whole*") (original emphasis).

For the combination of Takahashi and Hu as applied to the independent claims, the Board provides little more than the same conclusory assertions made for the dependent claims, except to further assert, "There would have been a reasonable expectation of success in doing so, as both references teach using information related to cache misses to determine when to power down a portion of a cache, and Hu additionally teaches turning off cache lines based on a comparison between an estimated power loss due to cache misses and an estimated power gain from turning off a portion of a cache." Appx36. The second clause in this sentence is merely a statement about the alleged teaching

---

349, Appx1910-1981, Appx2194-2221. By contrast, Patent Owner provided specific arguments and evidence demonstrating a lack of reasonable expectation of success for this combination. Appx277-279, Appx2148-2151.

of Hu, and does not bear upon the combination or provide any basis supporting a skilled artisan's expectation of success. Thus, the only basis upon which the Board reaches its conclusion is the assertion that the references both "teach using information related to cache misses to determine when to power down a portion of a cache." This is inadequate for at least two reasons. First, as discussed above in Section III regarding motivation to combine, Takahashi is demonstrably not directed to determining when to power down a portion of a cache based on information related to cache misses, but in fact is directed to determining when to *power up* a portion that has already been shut down. As such, the Board's finding lacks substantial evidence support on this basis alone. Second, the Board has resorted to a gross overgeneralization of the references to allege similarity where none exists. At best, Hu is concerned with the *energy* associated with a cache miss (*i.e.*, where an additional access to the L2 cache is required) while Takahashi is concerned with the *rate* of cache misses. Appx869 ("[A] cache-miss rate measuring circuit 140 measures the cache-miss rate during way access operation."); Appx889-890 (characterizing *L2Access* as "energy per cache access" and *leak* as "energy per cycle"). Notwithstanding a superficial similarity, these very different approaches, based on entirely different quantities having different values and even different units in a given system, demonstrate the *dissimilarity* of the references and the substantial

challenges a person of skill in the art would face in attempting to somehow combine their teachings. Accordingly, the Board's finding of reasonable expectation of success for the independent claims remains unsupported by substantial evidence.

Finally, for the combination of Takahashi with Hu and Cohen as applied to claims 4, 5, 15, and 16, the Board again fails to articulate a reasoned basis for its finding of reasonable expectation of success supported by substantial evidence. In this case, it relies on the Petitioner's arguments as set forth in its IPR Petition.[5] Specifically, the Board states, "Petitioner contends a person of ordinary skill in the art would have had a reasonable expectation of success in combining Takahashi, Hu, and Cohen due to their similarities," and deems this contention "adequate" to support its finding of reasonable expectation of success. Appx47 (citing Appx1975 and Appx863-864 (¶ 145)). The Petition, for its part, alleges that "[a] person of ordinary skill in the art would have had a reasonable expectation of success in combining Takahashi, Cohen, and Hu to arrive at the claimed inventions" because "the references are directed to addressing the same specific problem, using a similar technique and

[5] The combination of Takahashi with Hu and Cohen as applied to claims 4, 5, 15, and 16 appears to be the sole instance in which the Petitioner has actually alleged that a skilled artisan would have a reasonable expectation of success in making the proposed combination to reach the claimed invention, including in its IPR briefing and in its two expert declarations.

considerations." Appx1975.[6] The Board's reliance on such a conclusory and unsupported argument is insufficient to support its conclusion of obviousness. *In re Kahn*, 441 F.3d at 989. It also overlooks the substantial differences among the references, not least of which are the differences between Takahashi and Hu discussed above, which continue to impair the Board's finding as it relates to the further combination with Cohen. Accordingly, the Board's findings of obviousness are unsupported by substantial evidence and should be reversed.

CONCLUSION

The Board's Final Written Decisions should be reversed-in-part, vacated, and remanded.

Dated: April 21, 2023

Respectfully Submitted,

By: */s/ Kamran Vakili*

Kamran Vakili
IRELL & MANELLA LLP
840 Newport Center Dr., Ste. 400
Newport Beach CA 92660
Telephone: (949) 760-0991
Facsimile: (949) 760-5200

Attorney for Appellant
VLSI Technology LLC

[6] The exact language is repeated in Petitioner's expert declaration. Appx863-864 (¶ 145).

ADDENDUM

TABLE OF CONTENTS

Page(s)

Trials@uspto.gov
571-272-7822

Paper 30
Date: October 11, 2022

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,
Petitioner,

v.

VLSI TECHNOLOGY LLC,
Patent Owner.

IPR2018-01312
IPR2018-01661
Patent 8,020,014 B2

Before LYNNE E. PETTIGREW, BART A. GERSTENBLITH, and KIMBERLY McGRAW, *Administrative Patent Judges*.

PETTIGREW, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision on Remand
Determining All Challenged Claims Unpatentable
35 U.S.C. §§ 144, 318(a)

I. BACKGROUND

This Remand Decision is a final written decision on remand from the United States Court of Appeals for the Federal Circuit, which affirmed in part, reversed in part, and remanded the original Final Written Decisions in these two proceedings. *See Intel Corp. v. VLSI Tech. LLC*, 858 Fed. App'x 349 (Fed. Cir. 2021) (nonprecedential); IPR2018-01661, Paper 25 ("1661 Final Dec."); IPR2018-01312, Paper 26 ("1312 Final Dec.").

We have jurisdiction under 35 U.S.C. § 6, and we issue this Remand Decision under 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73. For the reasons discussed below, Petitioner has shown by a preponderance of the evidence that claims 1–5, 12–16, 18, and 20 of U.S. Patent No. 8,020,014 B2 (Ex. 1101,[1] "the '014 patent") are unpatentable.

A. Procedural History

In IPR2018-01661, Petitioner, Intel Corporation, filed a Petition requesting *inter partes* review of claims 1–5, 15, and 16 of the '014 patent. Paper 2 ("Pet.").[2] Patent Owner, VLSI Technology LLC, filed a Preliminary Response. Paper 7. After we instituted an *inter partes* review of the challenged claims (Paper 9, "1661 Inst. Dec."), Patent Owner filed a Patent Owner Response (Paper 13, "PO Resp."), Petitioner filed a Reply (Paper 15, "Pet. Reply"), and Patent Owner filed a Sur-reply (Paper 17, "PO Sur-reply").

---

[1] Exhibits having numbers beginning with "11" and "21" are from the record in IPR2018-01661; exhibits having numbers beginning with "10" and "20" are from the record in IPR2018-01312. When the same exhibit appears in both proceedings, we cite the exhibit from IPR2018-01661.

[2] Unless otherwise indicated, citations to papers refer to papers filed in IPR2018-01661.

In IPR2018-01312, Petitioner filed a Petition requesting *inter partes* review of claims 12–14, 18, and 20 of the '014 patent. IPR2018-01312, Paper 1 ("1312 Pet."). Patent Owner filed a Preliminary Response. IPR2018-01312, Paper 8. After we instituted an *inter partes* review of the challenged claims (IPR2018-01312, Paper 10, "1312 Inst. Dec."), Patent Owner filed a Patent Owner Response (IPR2018-01312, Paper 14, "1312 PO Resp."), Petitioner filed a Reply (IPR2018-01312, Paper 16, "1312 Pet. Reply"), and Patent Owner filed a Sur-reply (IPR2018-01312, Paper 18, "1312 PO Sur-reply").

A consolidated oral hearing for IPR2018-01312 and IPR2018-01661 was held on November 6, 2019, and the record of each proceeding includes a transcript of the hearing. Paper 24 ("Tr."); IPR2018-01312, Paper 25.

We issued Final Written Decisions in both cases, concluding that Petitioner had failed to demonstrate by a preponderance of the evidence that any of the challenged claims were unpatentable. *See* 1661 Final Dec. 27; 1312 Final Dec. 26. In IPR2018-01661, we determined that Petitioner had not shown that:

(1) claims 1–3 are unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi[3];

(2) claims 1–5, 15, and 16 are unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi and Hu[4]; and

[3] U.S. Patent No. 5,761,715, issued June 2, 1998 (Ex. 1103, "Takahashi").
[4] Zhigang Hu et al., *Let Caches Decay: Reducing Leakage Energy via Exploitation of Cache Generational Behavior*, ACM TRANSACTIONS ON COMPUTER SYSTEMS, May 2002, at 161–90 (Ex. 1104, "Hu").

(3) claims 4, 5, 15, and 16 are unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi, Hu, and Cohen.[5]

1661 Final Dec. 9–27.

In IPR2018-01312, we determined that Petitioner had not shown that:

(1) claims 12–14, 18, and 20 are unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi;

(2) claims 12–14, 18, and 20 are unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi and Hu;

(3) claim 20 is unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi and Gunther;[6] and

(4) claim 20 is unpatentable under 35 U.S.C. § 103(a) for obviousness over Takahashi, Hu, and Gunther.

1312 Final Dec. 9–26.

On appeal, the Federal Circuit affirmed our determination that Takahashi does not teach an "estimated power gain," as recited in independent claims 1 and 12. *Intel*, 858 Fed. App'x at 352–53. Petitioner's asserted obviousness grounds based on Takahashi alone and on Takahashi combined with Gunther depend on Petitioner's contention that Takahashi teaches the claimed "estimated power gain." Pet. 28–34; 1312 Pet. 32–36, 57. Therefore, in view of the Federal Circuit's affirmance regarding Takahashi, our findings and conclusions with respect to obviousness based on Takahashi and obviousness based on the combination of Takahashi and Gunther remain undisturbed. *See* 1661 Final Dec. 9–20; 1312 Final Dec. 9–19, 26.

---

[5] U.S. Patent No. 7,127,560 B2, issued Oct. 24, 2006 (Ex. 1105, "Cohen").
[6] U.S. Patent No. 5,781,783, issued July 14, 1998 (Ex. 1005, "Gunther").

The Federal Circuit reversed certain determinations we made as to the grounds involving the combination of Takahashi and Hu and remanded for further consideration of those grounds, including "whether the combination of Takahashi and Hu satisfies all of the claim limitations, and whether there was a motivation to combine Takahashi and Hu with a reasonable expectation of success." *Intel*, 858 Fed. App'x at 354–55. We did not authorize additional briefing or evidence on remand. Paper 28. Analyzing the full record in view of the Federal Circuit's directives, we address below the remaining grounds at issue based on (1) Takahashi and Hu, (2) Takahashi, Hu, and Cohen, and (3) Takahashi, Hu, and Gunther, summarized in the following tables:

IPR2018-01661:

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
1–5, 15, 16	103(a)[7]	Takahashi, Hu
4, 5, 15, 16	103(a)	Takahashi, Hu, Cohen

IPR2018-01312:

Claim(s) Challenged	35 U.S.C. §	Reference(s)/Basis
12–14, 18, 20	103(a)	Takahashi, Hu
20	103(a)	Takahashi, Hu, Gunther

B. The '014 Patent

The '014 patent, titled "Method for Power Reduction and a Device Having Power Reduction Capabilities," relates to power reduction of a cache

[7] The Leahy-Smith America Invents Act ("AIA"), Pub. L. No. 112-29, 125 Stat. 284, 287–88 (2011), amended 35 U.S.C. §§ 102 and 103, effective March 16, 2013. Because the '014 patent has an effective filing date before this date, we refer to the pre-AIA version of § 103. *See* Ex. 1101, code (22).

memory during a low power mode. *See* Ex. 1101, code (54), 1:7–10. In particular, the power management method disclosed in the '014 patent determines whether to power down a portion of a component, such as a cache memory, in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the portion of the component during a low power mode. *Id.* at code (57), 1:7–10, 4:31–39.

The '014 patent discloses that an estimated power gain resulting from powering down a portion of a component can be determined each time a request for powering down is received or in a periodical manner, in a random manner, in a pseudo-random manner, or in response to the occurrence of events (e.g., a certain battery level or a change in a clock frequency). *Id.* at 5:55–63. The power gain may be predefined and can be provided by the manufacturer of the device that includes the cache memory component. *Id.* at 5:64–67. An estimated power gain also may be based on the duration or other characteristics of the low power mode, which may be determined by monitoring operation of the device. *Id.* at 6:1–12.

The '014 patent also describes different ways of estimating a power loss resulting from powering down a portion of a component. For instance, an estimated power loss can be based on the amount of so-called "dirty information" stored in a cache memory. *Id.* at 6:36–37. When a cache implements a "write-back" policy, dirty information is information that has been modified and stored in cache memory but has not yet been written to high-level memory. *Id.* at 1:62–2:9. If a portion of the cache containing dirty information is powered down, power is consumed as the dirty information in cache memory first must be written to high-level memory, a process known as "flushing the cache." *Id.* at 1:63–67, 2:7–9, 3:43–47. The

power loss resulting from writing the dirty information to high-level memory as a result of powering down a cache portion may be estimated by counting or otherwise estimating the number of "dirty bits," also known as "sticky bits," which is a reflection of the amount of dirty information stored in the cache. *Id.* at 4:46–57, 6:32–41. In one embodiment, the '014 patent describes comparing an amount of dirty cache lines (i.e., portions of a cache containing dirty information) with "a threshold representative of the estimated power gain in order to determine the relationship between the power loss and power gain." *Id.* at 6:42–45; *see id.* at 4:64–67 (comparing amount of dirty information to "a power gating threshold TH" to determine whether to flush the cache and power down during a low power mode).

The '014 patent further describes generating cache statistics that can be used to estimate a power loss. *Id.* at 7:3–9. Cache statistics include cache hits, which occur when retrieval of information present in a cache is requested, and cache misses, which occur when retrieval of information that is not present in a cache is requested and has to be fetched from a higher-level memory module. *Id.* at 1:25–28.

Once the power management method described in the '014 patent determines an estimated power gain and an estimated power loss that would result from powering down a portion of a component, it determines whether to power down the portion of the component in response to a relationship between the estimated power gain and estimated power loss. *Id.* at 6:46–49. For example, if the estimated power gain is less than the estimated power loss, the cache memory may be shut down during a lower power mode. *Id.* at 6:49–51.

C. Illustrative Claims

Petitioner challenges claims 1–5, 12–16, 18, and 20 of the '014 patent. Claims 2–5 depend directly or indirectly from independent claim 1, and claims 13–16, 18, and 20 depend directly or indirectly from independent claim 12. Claims 1 and 12, reproduced below, are illustrative of the claimed subject matter:

> 1. A method for power reduction, the method comprises:
>
> [a] selectively providing power to at least a portion of a component of an integrated circuit during a low power mode; and
>
> [b] determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode.
>
> 12. A device having power reduction capabilities, the device comprises:
>
> [a] power switching circuitry adapted to selectively provide power to at least a portion of a component of the device during a low power mode, [b] having power management circuitry adapted to determine whether to power down at least the portion of the component during a low power mode [c] in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the lower power mode.

Ex. 1101, 7:27–35, 8:8–18 (modified to include lettering added by Petitioner).

D. Testimonial Evidence

In support of its unpatentability contentions in IPR2018-01661, Petitioner relies on a declaration by Anand Raghunathan, Ph.D. (Ex. 1102), filed with the Petition in that case, and a second declaration by Dr. Raghunathan (Ex. 1113), filed with Petitioner's Reply in that case. In

support of its Patent Owner Response in that case, Patent Owner relies on a declaration of Andrew Wolfe, Ph.D. (Ex. 2107).

In support of its unpatentability contentions in IPR2018-01312, Petitioner relies on a declaration by Dr. Raghunathan (Ex. 1002), filed with the Petition in that case, and a second declaration by Dr. Raghunathan (Ex. 1013), filed with Petitioner's Reply in that case. In support of its Patent Owner Response in that case, Patent Owner relies on a declaration of Dr. Wolfe (Ex. 2007).

Patent Owner cross-examined Dr. Raghunathan via deposition. *See* Ex. 2108. Petitioner cross-examined Dr. Wolfe via deposition. *See* Ex. 1112.

II. DISCUSSION

A. Legal Principles

A claim is unpatentable under 35 U.S.C. § 103(a) if the differences between the claimed subject matter and the prior art are "such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains." *KSR Int'l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations, including (1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and (4) when in evidence, objective indicia of non-obviousness. *Graham v. John Deere Co.*, 383 U.S. 1, 17–18 (1966).[8]

[8] With respect to the fourth *Graham* factor, the parties in these proceedings do not present arguments or evidence regarding objective indicia of non-obviousness. Therefore, the obviousness analysis below is based on the first three *Graham* factors.

To prevail on its challenges to Patent Owner's claims, Petitioner must demonstrate by a preponderance of the evidence that the claims are unpatentable. 35 U.S.C. § 316(e); 37 C.F.R. § 42.1(d). "In an [*inter partes* review], the petitioner has the burden from the onset to show with particularity why the patent it challenges is unpatentable." *Harmonic Inc. v. Avid Tech., Inc.*, 815 F.3d 1356, 1363 (Fed. Cir. 2016) (citing 35 U.S.C. § 312(a)(3) (requiring *inter partes* review petitions to identify "with particularity . . . the evidence that supports the grounds for the challenge to each claim")). This burden never shifts to Patent Owner. *See Dynamic Drinkware, LLC. v. Nat'l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015) (citing *Tech. Licensing Corp. v. Videotek, Inc.*, 545 F.3d 1316, 1326–27 (Fed. Cir. 2008)) (discussing the burden of proof in *inter partes* review).

An obviousness determination requires finding "both 'that a skilled artisan would have been motivated to combine the teachings of the prior art references to achieve the claimed invention, and that the skilled artisan would have had a reasonable expectation of success in doing so.'" *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1367–68 (Fed. Cir. 2016) (citation omitted); *see KSR*, 550 U.S. at 418 (holding that in an obviousness analysis, "it can be important to identify a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does"). Further, "[t]o satisfy its burden of proving obviousness, a petitioner cannot employ mere conclusory statements. The petitioner must instead articulate specific reasoning, based on evidence of record, to support the legal conclusion of obviousness." *In re Magnum Oil Tools Int'l, Ltd.*, 829 F.3d 1364, 1380 (Fed. Cir. 2016).

B. Level of Ordinary Skill in the Art

In the Final Written Decisions, we determined that "a person of ordinary skill in the art would have had at least a B.S. degree in computer science, computer engineering, or electrical engineering (or equivalent experience) and three years of experience with computer hardware and software power management design techniques." 1661 Final Dec. 8; 1312 Final Dec. 7–8. We see no reason to depart from this determination, which the parties did not challenge on appeal.

C. Claim Construction

We stated in the Final Written Decisions that terms in the challenged claims should be interpreted according to their broadest reasonable construction in light of the specification of the '014 patent because the Petitions were filed prior to November 13, 2018, and the '014 patent had not expired. 1661 Final Dec. 9 (citations omitted); 1312 Final Dec. 8–9 (citations omitted). We also stated that the parties had not proposed a construction for any term, and we determined it was unnecessary to construe any claim terms expressly to resolve the controversies before us. 1661 Final Dec. 9 (citations omitted); 1312 Final Dec. 9 (citations omitted). The Federal Circuit did not address the construction of any claim term. We continue to find it unnecessary to construe any claim terms expressly.

D. Overview of Asserted Prior Art

1. Takahashi

Takahashi relates to reducing power consumption in cache memories. Ex. 1103, 1:9–12. In the Background of the Invention section, Takahashi describes known prior art methods of reducing power consumption in a

cache, such as reducing the number of ways[9] to be executed in a cache memory. *Id.* at 1:35–57. Takahashi describes one prior art method in which "the number of ways to be executed in the cache memory [is] decreased, so that the cache-hit rate becomes low. Accordingly, the number of access operations to an external memory . . . is increased and the access operation speed becomes low." *Id.* at 1:58–63. "In other words, the performance of the information processing device becomes low because the operation time of the information processing device becomes long," causing increased power consumption of the information device. *Id.* at 1:63–2:1. In this method, "[t]he relationship between the low power consumption mode for a cache memory and the amount of actually reduced power consumption in the information processing device is determined based on the change of cache hit rate." *Id.* at 2:1–6. Because "the cache hit rate is also changed by application programs to be executed, . . . it is difficult to determine or know what the main cause to reduce the cache hit rate is." *Id.* at 2:6–9.

In view of this prior art, Takahashi concludes:

> In order to reliably reduce the power consumption of an information processing device in which the number of ways forming a cache memory system is decreased without reducing the cache hit rate, the number of ways to be executed in the cache memory system must be dynamically changed or switched while observing the change of cache hit rate.

Id. at 2:10–15. According to Takahashi, "there is no conventional cache memory system and no conventional information processing device having the cache memory system that satisfy this requirement." *Id.* at 2:16–21.

---

[9] Dr. Raghunathan explains that it was well known that "ways" and "lines" are portions of a cache. Ex. 1102 ¶¶ 24–25 (citing Ex. 1110, 397–99 (computer architecture textbook); Ex. 1101, 2:1–2, 2:54–57; Ex. 1103, 1:29–34).

Accordingly, an object of the invention described in Takahashi is to provide a cache memory and information processing device having the cache memory that are capable of adjusting the number of cache ways to be accessed "according to operation states of the information processing device, without decreasing the operation speed and the performance of the information processing device while decreasing the power consumption of the information processing device." *Id.* at 2:24–35.

Figure 1 of Takahashi is reproduced below:

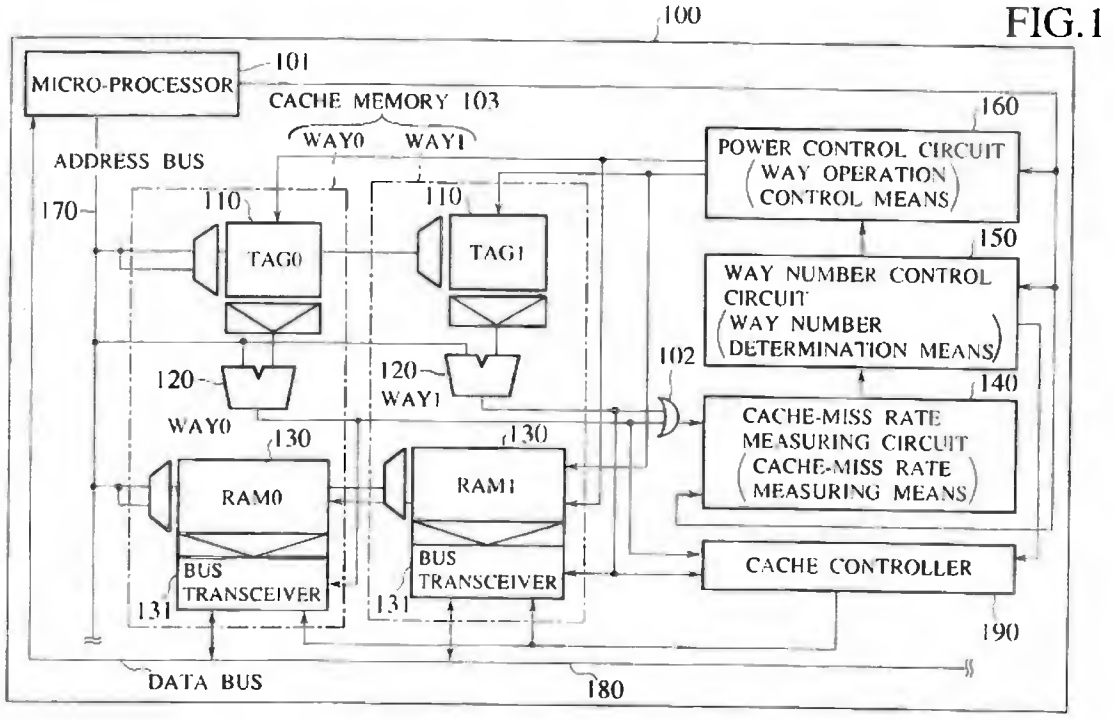

Id. at Fig. 1. As shown in Figure 1 above, Takahashi discloses an information processing device incorporating a cache memory that comprises (1) two cache ways (WAY0 and WAY1) in cache memory 103, (2) cache-miss rate measuring circuit 140, (3) way number control circuit 150, and (4) power control circuit 160. *Id.* at 2:36–47, 4:64–6:18, Fig. 1. The cache-miss rate measuring circuit measures the number of cache misses during a memory access operation mode. *Id.* at 2:48–50, 5:53–67,

6:59–7:2, Fig. 2. The cache-miss rate is transferred to the way number control circuit, which determines the optimum number of cache ways to be accessed based on the change of the cache-miss rate. *Id.* at 2:54–58, 5:67–6:6, Fig. 3. The way number control circuit transfers a control signal to the power control circuit, which supplies power to selected cache ways and does not supply power to other cache ways. *Id.* at 6:6–18, 8:27–28, 8:43–45. For example, for the cache memory shown in Figure 1, the power control circuit supplies power to either one cache way (i.e., WAY0 or WAY1) or both cache ways (i.e., WAY0 and WAY1). *See id.* at 6:19–21.

Figure 4 of Takahashi is reproduced below:

FIG.4

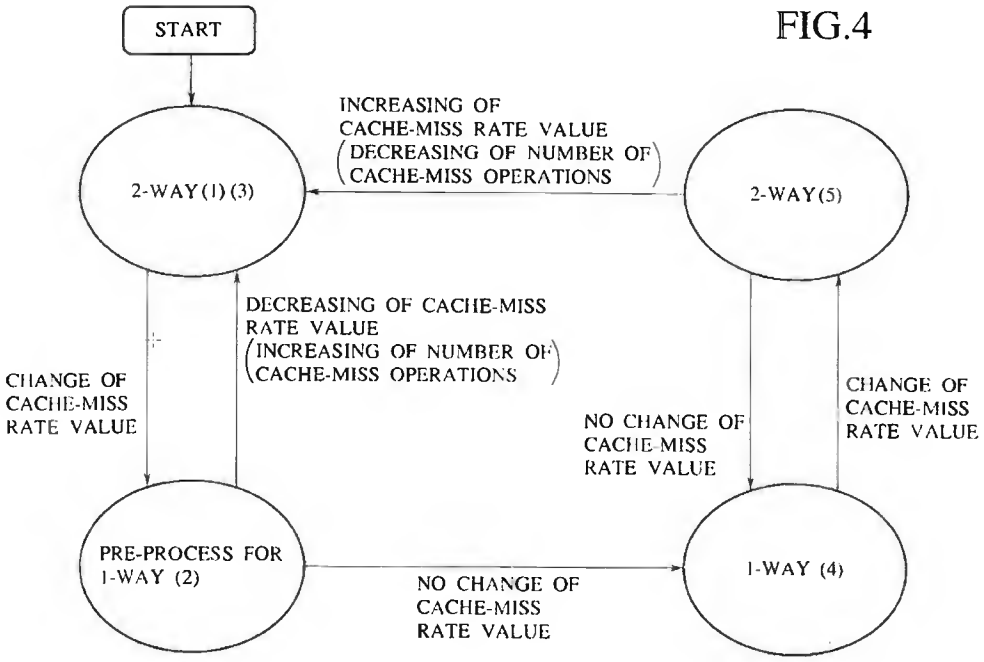

Figure 4, above, is a diagram showing operation flow of Takahashi's way number control circuit. *Id.* at 4:43–46, 9:37–40. Operation begins at (1) in "2-way" operation mode in which two cache ways are powered. *Id.* at 9:41–43. Operation then enters the "pre-process for 1-way" mode at (2), which "is executed before the operation flow is completely shifted to the one

way process in order to decrease the activated way number from two ways to one way." *Id.* at 9:50–54. In this mode, Takahashi's way control circuit determines "the optimum way number for the current process state." *Id.* at 9:54–55. When the cache-miss rate "does not exceed the predetermined value" (i.e., when the cache-miss rate does not increase), the way control circuit shifts to the "1-way" operation mode at (4) (i.e., it determines that only one cache way should be powered on). *Id.* at 9:55–59. On the other hand, if the cache-miss rate increases, the way control circuit returns to "2-way" mode at (3) (i.e., it determines that two cache ways should be powered on). *Id.* at 9:59–62.

2. Hu

Hu, an academic paper titled "Let Caches Decay: Reducing Leakage Energy via Exploitation of Cache Generational Behavior," examines methods for reducing leakage power within cache memories of CPUs. Ex. 1104, 161. In particular, Hu evaluates time-based cache decay policies in which a cache line is turned off if a preset number of cycles (i.e., a decay interval) have elapsed since the cache line was last accessed. *Id.* at 162. Hu explains that turning off cache lines during dead periods reduces leakage power dissipated by the cache lines storing data items that are no longer useful. *Id.* at 164. Hu recognizes, however, that turning off cache lines may introduce additional cache misses. *Id.* at 166. Therefore, a "basic premise" of Hu's evaluations "is to measure the static power saved by turning off portions of the cache, and then compare it to the extra dynamic power dissipated" when turning off cache lines "introduces additional L1 cache[10]

[10] The L1 and L2 caches are different levels of a memory hierarchy, with the L1 cache being the level closest to the CPU. *See* Ex. 1104, 166 (Table I).

misses [that] translate to extra L2 [cache] reads and sometimes also extra writebacks." *Id.* Hu states:

> We wish to turn off cache lines as often as possible in order to save leakage power. We balance this, however, against a desire to avoid increasing the miss rate of the L1 cache. Increasing the miss rate of the L1 cache has several power implications. First and most directly, it causes dynamic power dissipation due to an access to the L2 cache, and possible additional accesses down the memory hierarchy.

Id. at 167. Hu's cache decay policies attempt to "balance the potential for saving leakage energy (by turning lines off) against the potential for incurring extra level-two cache accesses (if we introduce extra misses by turning lines off prematurely)." *Id.* at 168.

A key energy metric in Hu's study is "normalized cache leakage energy," which "refers to a ratio of the energy of the L1 with cache decay policies versus the original L1 cache leakage energy." *Id.* at 166. The normalized cache leakage energy is a function of various weighted factors, including an "*L2Access:leak* ratio," which "relates dynamic energy due to an additional miss (or writeback) to a single clock cycle of static leakage energy in the L1 cache." *Id.* at 166–67. "The denominator of the *L2Access:leak* [ratio] relates to the leakage energy dissipated by the L1 data cache" during one clock cycle. *Id.* at 168. This static leakage energy is the power that would be saved by turning off a portion of the L1 cache. *See id.* at 166. The numerator of the *L2Access:leak* ratio represents the amount of dynamic power dissipated by an access (read or writeback) to the L2 cache, and possibly to other levels of the memory hierarchy, due to an L1 cache miss that would result from turning off portions of the L1 cache. *Id.* at 166, 167. "Multiplying [the *L2Access:leak* ratio] by the number of extra L2 accesses induced by cache decay gives the dynamic cost induced" when a

cache line is turned off. *Id.* at 167; *see id.* at 166–67 (identifying (*L2Access:leak*)(*extraL2Accesses*) as one of three terms in Hu's "normalized cache leakage energy").

Based on industry data and recent work estimating the dynamic energy per L2 cache access and the leakage energy dissipated by the L1 data cache, Hu estimates *L2Access* to be in the range of 3 to 5 nJ per L2 access for a 1 MB L2 cache and *leak* to be an average of 0.45 nJ per clock cycle for a 32 KB L1 cache. *Id.* at 167–68. Combining the two, Hu estimates an *L2Access:leak* ratio of 8.9 for devices using caches of those sizes. *Id.* at 168. Hu also explains that *L2Access:leak* ratio estimates will vary with design style and fabrication technology and that leakage energy is expected to increase in the future. *Id.* To account for all these factors, Hu evaluates time-based cache decay policies by determining the normalized cache leakage energy across a set of benchmark applications for different values of the *L2Access:leak* ratio, including ratios of 5, 10, 20, and 100. *Id.* at 165–66, 173–76, Figs. 8, 9.

3. Cohen

Cohen describes a method of powering down sections of a cache to minimize power consumption, while not impacting performance when high performance is required, and to minimize leakage power from a cache when full performance is not required. Ex. 1105, 1:29–36. Cohen describes the operation of a writeback cache, in which data written to the cache is not written to main memory until a valid cache line is needed for new data. *Id.* at 2:61–65. Data contained in a writeback cache that has not yet been stored elsewhere is referred to as "dirty" data. *Id.* at 2:65–67. Cohen describes various ways of storing data from dirty cache lines in memory before a portion of a cache can be powered down. *Id.* at 3:65–4:49. Cohen

also describes a busy bit counter that counts the number of dirty bits in a cache. *Id.* at 4:64–5:4.

E. Asserted Obviousness over Takahashi and Hu

Petitioner contends that claims 1–5, 12–16, 18, and 20 are unpatentable under 35 U.S.C. § 103(a) for obviousness over the combined teachings of Takahashi and Hu. Pet. 36–56; 1312 Pet. 43–56. Petitioner alleges that the combination of Takahashi and Hu teaches or suggests the subject matter of the challenged claims, and Petitioner provides reasons why a person of ordinary skill in the art allegedly would have combined the references. Pet. 36–56; 1312 Pet. 43–56; *see* Ex. 1102 ¶¶ 88–131; Ex. 1002 ¶¶ 107–135. In response, Patent Owner argues that Petitioner fails to show that the combination of Takahashi and Hu teaches or suggests every limitation of independent claims 1 and 12 and fails to show that a person of ordinary skill in the art would have combined the references in the manner asserted. PO Resp. 20–33; 1312 PO Resp. 21–34; *see* Ex. 2107 ¶¶ 46–61; Ex. 2007 ¶¶ 48–62. Patent Owner also argues that Petitioner fails to show that the combination of Takahashi and Hu teaches the subject matter of dependent claims 4, 5, 15, 16, 18, and 20. PO Resp. 38–40; 1312 PO Resp. 19–21, 39; *see* Ex. 2107 ¶¶ 59–61; Ex. 2007 ¶¶ 45–47, 61. Petitioner replies to Patent Owner's arguments (Pet. Reply 14–20; 1312 Pet. Reply 15–20), and Patent Owner responds to the reply arguments (PO Sur-reply 14–24; 1312 PO Sur-reply 15–24).

For the reasons explained below, Petitioner has established by a preponderance of the evidence that claims 1–5, 12–16, 18, and 20 of the '014 patent would have been obvious over the combined teachings of Takahashi and Hu.

1. Claim 1

a. Preamble and Limitation 1[a]

Petitioner contends that Takahashi teaches a "method for power reduction," as recited in the preamble of claim 1. Pet. 25; *see id.* at 36 (referring back to obviousness analysis based on Takahashi alone). Specifically, Petitioner contends that Takahashi discloses "a cache memory and an information processing device having the cache memory" that "is capable of adjusting . . . the number of ways to be accessed in the cache memory . . . without decreasing the operation speed and the performance of information processing device while decreasing the power consumption of the information processing device." *Id.* at 25 (quoting Ex. 1103, 2:27–35) (citing Ex. 1103, 2:23–4:23, 4:31–42, 4:64–5:13, Figs. 1, 2, 3; Ex. 1102 ¶ 65). Petitioner also asserts that Takahashi discloses a power control circuit that receives a signal from a way number control circuit and shifts to the "one way process" operation, at which time the power supply to a specific cache way is turned off (i.e., "selectively providing power to at least a portion of a component"). *Id.* at 25–26 (citing, e.g., Ex. 1103, 6:10–18, 8:27–28, 8:43–45, 9:37–10:33, Fig. 4; Ex. 1102 ¶ 67). Petitioner further contends that Takahashi teaches that its power saving techniques, including powering down a cache way, are employed in a "low power consumption mode" (i.e., a "low power mode"). *Id.* at 27–28 (citing, e.g., Ex. 1103, 1:58–2:6; Ex. 1102 ¶ 69).

Patent Owner does not specifically dispute that Takahashi teaches the claim 1 preamble and limitation 1[a]. *See* PO Resp. Based on the complete record, and for the reasons explained in the Petition, Petitioner has

demonstrated sufficiently that Takahashi teaches the subject matter of claim 1's preamble and limitation 1[a].[11]

b. Limitation 1[b] and Motivation to Combine Takahashi and Hu

Claim 1 further recites "determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode." Ex. 1101, 7:31–35 (limitation 1[b]). The dispute between the parties focuses on whether the combination of Takahashi and Hu teaches this limitation and whether a person of ordinary skill in the art would have been motivated to combine the references as asserted. Below, we consider the parties' arguments based on the complete record and in view of (1) the Federal Circuit's determination that Hu teaches using the *L2Access:leak* ratio to determine whether to power down a portion of a cache, and (2) the Federal Circuit's instruction that we give due consideration to an argument that Takahashi teaches "determining whether to power down," which could be combined with Hu's teaching of "a relationship between an estimated power gain and an estimated power loss." *See Intel*, 858 Fed. App'x at 354–55.

Petitioner first relies on Takahashi for teaching limitation 1[b]. Pet. 37 (referring back to obviousness analysis based on Takahashi alone); *see id.* at 28–34. In particular, Petitioner contends that Takahashi discloses a way number control circuit that determines, during a low power mode, whether to power down a cache way based on a comparison of the cache-miss rate (representing power loss) to a predetermined value (representing

[11] We need not decide whether the preamble is limiting because we find that Petitioner has shown that Takahashi teaches it.

power gain). *Id.* at 29–30 (citing Ex. 1103, 6:3–18, 8:46–50, 9:37–10:33); *see* Ex. 1103, 9:50–59 ("In the pre-process operation mode for one way process, . . . when the cache-miss rate does not exceed the predetermined value, the process flow of the cache memory 103 is shifted to the one way process operation . . . ."). The Federal Circuit affirmed our determination in the Final Written Decision that Takahashi's predetermined value is a previous measurement of the cache-miss rate and not an "estimated power gain" as recited in limitation 1[b]. *See Intel*, 858 Fed. App'x at 352–53; 1661 Final Dec. 17–19 (citing Ex. 1103, code (57), 3:23–31, 9:55–61; Ex. 1112, 131:9–132:5). Our Final Written Decision did not address whether Takahashi teaches the other aspects of limitation 1[b]. *See* 1661 Final Dec. 20, 25.

Petitioner additionally relies on Hu for teaching limitation 1[b]. Pet. 37–42. As explained above, Hu teaches that turning off a portion of a cache saves static leakage energy, resulting in a power gain. Ex. 1104, 166–68. At the same time, Hu teaches that turning off a portion of a cache may introduce additional cache misses, which cause additional accesses to another level of memory (e.g., Hu's L2 cache), resulting in dynamic power consumption (i.e., a power loss). *Id.* Hu explains that a "key aspect of [its cache decay] policies is the desire to balance the potential for saving leakage energy (by turning lines off) against the potential for incurring extra level-two cache accesses (if we introduce extra misses by turning lines off prematurely)." *Id.* at 168. Hu also explicitly discloses comparing the estimated power saved by turning off a cache line to the estimated power dissipated when turning off the cache line causes additional cache misses and additional L2 accesses. *Id.* at 166.

In view of these disclosures, Petitioner contends that Hu teaches (1) "measur[ing] the estimated power gain (the static/leakage power saved by turning off portions of the cache)," (2) "estimat[ing] the power loss (extra dynamic power) from powering down the individual lines of cache," and (3) "compar[ing] the two to determine whether to power down a portion of the cache." Pet. 37. According to Petitioner, these steps are reflected in Hu's *L2Access:leak* ratio, which "'relates dynamic energy due to an additional miss (or writeback) to a single clock cycle of static leakage energy in the L1 cache.'" *Id.* at 37–38 (quoting Ex. 1104, 167) (citing Ex. 1102 ¶ 92). As Petitioner notes, Hu evaluates decay policies using the normalized cache leakage energy, focusing specifically on values of the *L2Access:leak* ratio in its analysis. *Id.* at 38 n.5 (citing Ex. 1104, 167, 172; Ex. 1102 ¶¶ 93–94).

First, Petitioner contends that the denominator of the *L2Access:leak* ratio represents "an estimated power gain . . . resulting from powering down" cache lines in a low power mode. Pet. 38. Specifically, Petitioner argues that Hu discloses that the denominator of Hu's *L2Access:leak* ratio "relates to the leakage energy dissipated by the L1 data cache." *Id.* (citing Ex. 1104, 168). Based on data obtained from several sources, Hu teaches that an estimated average leakage energy of 0.45 nJ per clock cycle will be saved when a 32 KB L1 cache is powered down. Ex. 1104, 168; *see* Pet. 38–39.

Second, Petitioner contends that the numerator of the *L2Access:leak* ratio represents "an estimated power loss resulting from powering down" cache lines in a low power mode. Pet. 39. Specifically, Petitioner argues that Hu discloses that the numerator of Hu's *L2Access:leak* ratio relates to the dynamic energy used per L2 cache access due to an L1 cache miss. *Id.*

(citing Ex. 1104, 167). Hu cites recent work that estimates dynamic energy per L2 cache access in the range of 3 to 5 nJ per access for L2 caches that are 1 MB in size and cites other studies suggesting that additional L2 accesses consume up to 10 nJ. Ex. 1104, 167–68; *see* Pet. 39.

Third, Petitioner contends that "Hu uses the '*L2Access:leak* ratio' to determine whether and when to power down a cache line based on the relationship between an estimated power loss (in the numerator) and an estimated power gain (in the denominator)." Pet. 39. Petitioner cites Hu's Figure 9, which plots four curves of normalized cache leakage energy, corresponding to *L2Access:leak* ratios of 5, 10, 20, and 100, for increasing decay intervals (i.e., the number of cycles since the last cache access). *Id.* at 40 (citing Ex. 1104, 174, 176, Fig. 9; Ex. 1102 ¶ 98). Petitioner contends that a person of ordinary skill in the art "would have understood that the 'minimum' value of each curve in Hu's Figure 9 represents the optimal idle period (*i.e.*, how long the system should wait) before a cache line is powered down because the minimum value of each curve represents the lowest normalized leakage energy consumed." *Id.* at 40–41. Accordingly, Petitioner contends, a person of ordinary skill in the art "reviewing Hu's Figure 9 would have understood that the higher the *L2Access:leak* ratio (*i.e.*, the higher the power loss per cache miss), the longer the system should wait before powering down the cache in order to reduce the likelihood of additional power-consuming cache misses." *Id.* at 41. Therefore, according to Petitioner, "Hu uses the *L2Access:leak* ratio to determine the relationship between estimated power loss (numerator) and estimated power gain (denominator), and determines when to power down a cache line based on that relationship." *Id.* (citing Ex. 1102 ¶ 100).

Petitioner further contends that a person of ordinary skill in the art would have been motivated to combine Hu's disclosure of powering down based on an estimated power gain and power loss with Takahashi's teaching of powering down a portion of a cache only when the cache-miss rate does not exceed a predetermined value.[12] *See* Pet. 43–44 (citing Ex. 1102 ¶ 104; Ex. 1103, 9:41–62, 10:23–30; Ex. 1104, 166, 167). In the combination, the information processing device would power down a portion of a cache in response to a comparison between an estimated power gain (cache leakage power) and an estimated power loss (additional power consumed by memory accesses resulting from cache misses), as taught in Hu, rather than based on a change in the cache-miss rate, as taught in Takahashi. *See id.* at 44 (citing Ex. 1102 ¶ 104). Petitioner contends that combining the references in this way would "further Takahashi's goal of using information relating to the number of cache hits and misses to more efficiently decrease the power consumption of a device." *Id.* (citing Ex. 1102 ¶ 104).

Petitioner articulates several reasons why a person of ordinary skill in the art would have been motivated to combine the teachings of Takahashi and Hu in this manner. *Id.* at 42–46. First, Petitioner contends that Takahashi and Hu address the same problem of reducing power consumption in a cache memory. *Id.* at 42 (citing Ex. 1103, 2:23–35; Ex. 1104, 161 (Abstract); Ex. 1102 ¶ 102). Petitioner also contends that Takahashi and Hu solve this problem using similar techniques. *Id.* In particular, Petitioner asserts that Takahashi discloses monitoring changes in cache-miss rates to determine when to apply the well-known technique of

[12] As noted above, Takahashi's "predetermined value" is a previous measurement of the cache-miss rate. *See Intel*, 858 Fed. App'x at 352–53; 1661 Final Dec. 17–19 (citations omitted).

reducing power consumption by selectively powering down portions of a cache, and Hu similarly discloses measuring the energy dissipated as a result of additional cache misses to determine when to power down cache lines. *Id.* at 42–43 (citing Ex. 1101, 2:21–22, 2:54–60; Ex. 1103, 6:3–7, 8:12–50, 9:55–59; Ex. 1104, 164, 166–68; Ex. 1102 ¶ 103).

Additionally, Petitioner contends that combining Hu's teachings regarding powering down based on estimated power loss and estimated power gain with Takahashi's information processing device would have been a "routine modification of Takahashi, leading to expected results." *Id.* at 44 (citing *KSR*, 550 U.S. at 417). First, Petitioner contends an ordinarily skilled artisan easily would have modified Takahashi's cache-miss rate to reflect the estimated power loss resulting from powering down a cache way, as Hu discloses several studies that effectively measured estimated power losses based on the dynamic energy per L2 cache access. *Id.* at 45 (citing Ex. 1104, 167). Second, Petitioner contends an ordinarily skilled artisan easily would have modified Takahashi to estimate the power gain resulting from powering down a cache way, which Hu discloses can be estimated based on the leakage energy of the cache. *See id.* (citing Ex. 1104, 168). If Takahashi is modified to use estimates for power loss and power gain, as taught in Hu, Petitioner contends that "the resulting device would have the expected result of powering down the cache way when the estimated power loss from powering down a cache way does not exceed the estimated power gain." *Id.* (citing Ex. 1102 ¶ 105). Petitioner concludes that a person of ordinary skill in the art would have been motivated to combine the teachings of Takahashi and Hu

> to achieve the common goal of power reduction by efficiently powering down portions of a cache (*i.e.*, a cache way or a cache

> line) based on the relationship between an estimated power gain (*i.e.*, leakage energy saved by powering down the cache way or cache line) and estimated power loss (*i.e.*, energy costs resulting from the additional cache misses caused by powering down the cache way or cache line).

Id. at 45–46 (citing Ex. 1103, 2:1–5, 9:55–58; Ex. 1104, 167; Ex. 1101 ¶ 106).

In response to Petitioner's contentions, Patent Owner argues that Hu does not teach or suggest a "relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode." PO Resp. 20–25. First, Patent Owner contends that Hu does not teach "an estimated power gain" and "an estimated power loss" because Hu explicitly disavows estimating values for "static/leakage power" and "extra dynamic power." *Id.* at 20 (citing Pet. 37). Because Hu states "it is difficult to nail down specific values for evaluation purposes" and instead "focus[es] . . . on ratios of values," particularly the *L2Access:leak* ratio, Patent Owner argues that Hu teaches away from estimating values for leakage power (corresponding to power gain) and dynamic power (corresponding to power loss). *Id.* at 20–22 (citing, e.g., Ex. 1104, 166–67; *DePuy Spine Inc. v. Medtronic Sofamor Danek*, 567 F.3d 1314, 1326–27 (Fed. Cir. 2009); *In re Kahn*, 441 F.3d 977, 990 (Fed. Cir. 2006); *In re Fulton*, 391 F.3d 1195, 1201 (Fed. Cir. 2004)). Second, Patent Owner argues that Hu's *L2Access:leak* ratio does not involve estimated values "*resulting from powering down* the at least portion of the component," but that Hu instead selects values of the *L2Access:leak* ratio as inputs to its simulation. *Id.* at 22–25 (emphasis modified); PO Sur-reply 15–16.

Based on the arguments and evidence, we agree with Petitioner that Hu describes a "relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component" as claimed. First, we do not agree with Patent Owner's argument that Hu teaches away from estimating leakage power and extra dynamic power. To the contrary, Hu explicitly describes using estimates for the L1 static leakage energy and L2 access energy. Ex. 1104, 167 (referring to an "energy *estimate*" based on "recent work that *estimates* dynamic energy per L2 access in the range of 3 to 5 nJ" (emphasis added)), 168 (referring to 0.45 nJ "static leakage per cycle *estimate*" obtained from literature review (emphasis added)).

We also agree with Petitioner that Hu teaches an estimated power gain and an estimated power loss "resulting from powering down." Hu explains that the leakage energy in the denominator of the *L2Access:leak* ratio reflects the estimated static leakage power that would be saved by turning off a portion of the L1 cache (i.e., "an estimated power gain . . . resulting from powering down" a portion of the L1 cache). *Id.* at 166 (referring to leakage energy "saved by turning off portions of the cache"), 167 ("We wish to turn off cache lines as often as possible in order to save leakage power."). Hu also makes clear that the extra dynamic power in the numerator of the *L2Access:leak* ratio reflects the estimated amount of dynamic energy dissipated per L2 cache access due to an L1 cache miss that would result from turning off a portion of the L1 cache (i.e., "an estimated power loss resulting from powering down" a portion of the L1 cache). *Id.* at 167 ("The predominant effect to model is the amount of dynamic power dissipated in the level-two cache and beyond, due to the level-one cache miss [resulting from turning off cache lines]."), 168 (describing "the potential for incurring

extra level-two cache accesses (if we introduce extra misses by turning lines off prematurely)"). Thus, contrary to Patent Owner's argument, Hu teaches that the values used in the *L2Access:leak* ratio are an estimated power gain (saved L1 cache leakage energy) and an estimated power loss (additional energy used to access the L2 cache due an L1 cache miss) that would result from powering down a portion of the L1 cache. *See id.* at 166–76.

Patent Owner further argues that Hu does not describe "determining whether to power down" portions of a cache in response to values in the *L2Access:leak* ratio. PO Resp. 21 n.5, 24; PO Sur-reply 16–18. The Federal Circuit, however, found that "Hu teaches using the *L2Access:leak* ratio to estimate the optimal decay interval, which is in turn used to determine when to power down a cache." *Intel*, 858 Fed. App'x at 355. Accordingly, the court concluded that "Hu discloses using the *L2Access:leak* ratio to determine whether to power down a cache (or portion thereof)." *Id.*

Finally, Patent Owner argues that a person of ordinary skill in the art would not have been motivated to combine Takahashi and Hu. PO Resp. 25–33; PO Sur-reply 19–23. Primarily, Patent Owner contends that modifying Takahashi to incorporate Hu's teachings would frustrate Takahashi's stated purpose and change its principle of operation. PO Resp. 25–28 (citing, e.g., *Plas-Pak Indus., Inc. v. Sulzer Mixpac AG*, 600 Fed. App'x 755, 760 (Fed. Cir. 2015)); PO Sur-reply 20–21. According to Patent Owner, Takahashi's basic principle of operation involves powering down a cache way before measuring a change in the cache-miss rate and then, based on the change in the cache-miss rate, determining whether to keep that cache way powered down or to power it back up. PO Resp. 27 (citing Ex. 1103, code (57), 9:41–62; Ex. 2107 ¶ 52; Ex. 2108, 137:19–

139:19).[13] Patent Owner argues that modifying Takahashi with Hu as proposed by Petitioner would require "flipping the order" to "perform[] some estimate of the relevant criteria *before* making a determination of whether to power down in response to that analysis." *Id.* (citing Ex. 2107 ¶ 52). In its Reply, Petitioner contends that the proposed combination of Takahashi and Hu would not require such a modification because Takahashi determines the cache-miss rate based on the assumption that one cache way is powered down without actually powering down the cache way. Pet. Reply 17–18.

Each party supports its position by citing portions of Takahashi describing its pre-process operation. Patent Owner argues that the following passage shows that Takahashi powers down one cache way in the pre-process mode before measuring the cache-miss rate: "First, we assume that the number of ways is decreased from two ways to one way. In this case, way 1 is eliminated from the access operation. In other words, the power source is disconnected from the way 1." Ex. 1103, 8:61–64; *see* PO Resp. 7 (addressing obviousness based on Takahashi alone); PO Sur-reply 3–4

[13] In opposing Petitioner's contention that claim 1 would have been obvious over Takahashi alone, Patent Owner argues that Takahashi shuts down a cache way before measuring the change in cache-miss rate and, therefore, does not determine whether to power down a cache way "in response to" a relationship between an estimated power gain and an estimated power loss. PO Resp. 4–10. Patent Owner's argument that modifying Takahashi in view of Hu would change Takahashi's principle of operation refers back to its analysis of the obviousness challenge based on Takahashi alone. *Id.* at 27 ("However, as already discussed above, Takahashi's basic principle of operation is to power down a cache way first to induce a change in the cache-miss rate (which is then measured to make a further determination of whether to keep the way powered down or to power it back up again)."); *see* PO Sur-reply 19.

(same). Patent Owner also cites the following disclosure from Takahashi: "In the pre-process operation mode . . . if the cache-miss rate is increased, the number of the activated ways is changed from one way to two ways." Ex. 1103, 9:50–61; *see* PO Sur-reply 4. In Patent Owner's view, this shows that only one cache way is powered on during the pre-process operation, i.e., that Takahashi powers down the other cache way before measuring the change in the cache-miss rate. PO Sur-reply 4–5.

Patent Owner contends that its understanding of Takahashi's pre-process operation is consistent with Takahashi's Background of the Invention section, which identifies a need in the art to check the change of the cache hit rate while dynamically switching the number of ways to be accessed. *Id.* at 5 (citing Ex. 1103, 2:16–21). In addition, Patent Owner points out that Takahashi makes several references to measuring the cache-miss rate, allegedly showing that Takahashi measures the cache-miss rate after powering down a cache way, rather than simulating a cache-miss rate based on an assumed shutdown, as Petitioner asserts. *Id.* at 5–6 (citing Ex. 1103, 2:38–46, 3:23–31).

Patent Owner further contends that both parties' experts agree that Takahashi already has one way powered down in the pre-process mode. *Id.* at 6–7. For example, Petitioner's expert, Dr. Raghunathan, testified that Takahashi "states the power source is disconnected from way one." Ex. 2108, 137:9–25 (referring to Takahashi's statement that "the power source is disconnected from the way 1" (Ex. 1103, 8:63–64)); *see* PO Sur-reply 6–7. Dr. Raghunathan also testified that Takahashi describes changing from one cache way in the pre-process mode back to two cache ways. Ex. 2108, 139:24–140:9; *see* PO Sur-reply 7. Patent Owner's expert, Dr. Wolfe, agreed that the pre-process operation occurs after the number of

cache ways has been reduced to one. Ex. 1112, 122:1–11; *see* PO Sur-reply 7.

Petitioner argues in its Reply that Takahashi determines the cache-miss rate based on the assumption that one cache way is powered down without actually powering down the cache way. Pet. Reply 17–18. Petitioner asserts that, in the pre-process mode, "Takahashi estimates the impact of powering down a portion of the cache before powering it down so as not to increase the cache-miss rate." *Id.* at 6 (citing Ex. 1113 ¶ 7). Petitioner emphasizes Takahashi's disclosure of "a pre-process operation which is executed *before the process to decrease the number of ways*." Ex. 1101, 8:53–55 (emphasis added); *see* Pet. Reply 6. Petitioner also cites Takahashi's statement that "[i]n the pre-process operation mode for one way process, the pre-process is executed *before the operation flow is completely shifted to the one way process in order to decrease the activated way number from two ways to one way*." Ex. 1101, 9:50–55 (emphasis added); *see* Pet. Reply 6. Based on these passages, Petitioner concludes that "the pre-process step occurs before the number of ways is decreased, *i.e.*, before a portion of the cache is powered down." Pet. Reply 6 (citing Ex. 1113 ¶ 8).

Takahashi also discloses that "[t]he pre-process operation must be executed in order to avoid the increasing of the cache-miss rate temporally when required data items are stored in the way that is eliminated from the access operation without executing of the pre-process operation." Ex. 1103, 8:56–60; *see id.* at 9:21–22 ("It can be avoided by the pre-process operations described above that the cache-miss rate temporally increased."). Petitioner interprets this to mean that the pre-process operation must be executed without increasing the cache-miss rate and, therefore, Takahashi must not

contemplate powering down a portion of the cache during the pre-process operation. Pet. Reply 6–7.

Moreover, Petitioner contends, Takahashi's description of the pre-process operation is consistent with the cache being fully powered on. *Id.* at 7. In a first pre-process step, Takahashi explains that when a cache miss occurs, the data item must be stored in way 0, which Petitioner contends would be consistent either with Takahashi simulating the effect of way 1 being powered down or with way 1 actually being powered down. *Id.* (citing Ex. 1003, 9:1–2; Ex. 1113 ¶ 9). In the second step, there is a cache hit on data during a read operation from way 1, and, in the third step, there is a write operation to way 1. *Id.* at 7–8 (citing Ex. 1103, 9:14–20). With support from Dr. Raghunathan, Petitioner asserts that the read and write operations in the last two steps can occur only if way 1 is powered on. *Id.* (citing Ex. 1113 ¶¶ 10–11). Petitioner also cites the testimony of Patent Owner's expert, Dr. Wolfe, who agreed that a processor can read data from and write data to a cache way only if the cache way is powered on. *Id.* at 8 (citing Ex. 1112, 48:4–7, 45:12–16). Thus, Petitioner contends that Takahashi's pre-process mode "describes a method for estimating the impact on the cache-miss rate, *if* a portion of the cache *were* powered down." *Id.* at 8–9 (emphasis modified) (citing Ex. 1113 ¶ 13).

Based on the parties' arguments and evidence, we disagree with Patent Owner that combining Hu's teachings with those of Takahashi would alter Takahashi's principle of operation in a manner that would render the combination nonobvious. As set forth above, both parties' positions regarding the pre-process operation have support in Takahashi's disclosure, which unfortunately is not a model of clarity. The cited expert testimony does not resolve the dispute because it focuses on isolated passages of

Takahashi. Reading Takahashi in its entirety, we find it ambiguous on this point, i.e., whether it fully powers down a cache way before measuring a change in the cache-miss rate. Patent Owner's argument as to Takahashi's principle of operation would require a person of ordinary skill in the art to have understood that Takahashi unambiguously teaches powering down a cache way before determining a change in the cache-miss rate. Because we find that Takahashi does not unambiguously teach powering down before determining a change in the cache-miss rate, we do not agree with Patent Owner's argument.

Moreover, regardless of whether Takahashi powers down a cache way before or after measuring the cache-miss rate, Takahashi's pre-process operation determines the optimal number of cache ways based on changes in the cache-miss rate resulting from powering down a cache way. Ex. 1103, 2:54–58, 5:67–6:6. This determination achieves Takahashi's stated purpose of reducing power consumption without negatively impacting the information processing device's performance. *See* Ex. 1103, code (54), 2:24–35; Pet. 44. Thus, we do not agree with Patent Owner that modifying Takahashi to incorporate Hu's teaching of comparing estimated power loss and estimated power gain to determine when to power down a cache portion would frustrate this stated purpose.[14]

Next, Patent Owner argues that combining Takahashi and Hu as Petitioner proposes would require making Hu's decay interval depend on the *L2Access:leak* ratio or modifying Hu so that cache lines are powered down based on values of the *L2Access:leak* ratio instead of the elapsing of a decay

[14] We also do not agree with Patent Owner's unsupported assertion in its Sur-reply that Takahashi's stated purpose is to avoid increasing the cache-miss rate. *See* PO Sur-reply 21.

interval. PO Resp. 28–29; PO Sur-reply 19–20. Patent Owner is incorrect that such a modification to Hu would be needed. The Federal Circuit explicitly found that Hu teaches using the *L2Access:leak* ratio to determine whether to power down a portion of a cache. *See Intel*, 858 Fed. App'x at 355. Thus, Hu teaches determining whether to power down a cache line in response to a relationship between an estimated power gain (e.g., the static leakage power that would be saved by turning off an L1 cache line) and an estimated power loss (e.g., dynamic energy consumption by the L2 cache due to L1 cache misses resulting from turning off an L1 cache line). *See id.*; Ex. 1104, 167–68.

Patent Owner further argues that the combination of Takahashi and Hu would require a modification to provide some way of measuring or estimating *L2Access* and *leak* values for the system in question and there would be no expectation of success in modifying Takahashi and Hu to perform estimates of power gain and power loss. PO Resp. 30–31; PO Sur-reply 20. We do not agree. Claim 1 does not require measuring or performing estimates of power gain and power loss as part of the claimed method—it only requires determining whether to power down in response to a relationship between an estimated power gain and an estimated power loss. *See* Ex. 1101, 7:31–35. Furthermore, as the Federal Circuit found, Hu demonstrates that estimates of power gain and power loss can be based on empirically measured values for *L2Access* and *leak* for a given memory structure. *See Intel*, 858 Fed. App'x at 355; *see also* Ex. 1104, 167 (citing studies estimating power losses based on the dynamic energy per L2 cache access), 168 (estimating L1 cache leakage energy per cycle based on data obtained from several sources). Thus, the teachings of Takahashi and Hu

can be combined to teach limitation 1[b] without any modifications to perform estimates of power gain and power loss.

In summary, based on the complete record, we find that Petitioner has shown persuasively that the combination of Takahashi and Hu teaches "determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode." *See* Pet. 37–46. Takahashi teaches determining the number of cache ways that should be powered on based on a change in the cache-miss rate that would result from powering down a cache way in a low power consumption mode. *See* Ex. 1103, 9:50–62. Hu teaches comparing an estimated power gain (e.g., static L1 cache leakage power saved by turning off a cache line) with an estimated power loss due to additional cache misses (e.g., extra dynamic power dissipated by accessing the L2 cache) that would result from powering down a cache line and determining whether to power down a cache line by balancing the two. *See* Ex. 1104, 166–69, 173–76. Thus, when Takahashi is modified to incorporate Hu's comparison of static cache leakage power with extra dynamic power dissipated due to additional cache misses, the combination teaches "determining whether to power down [a cache way] in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the [cache way] during [a] low power mode," as recited in limitation 1[b]. *See* Pet. 42–46. Because we find that Hu teaches determining whether to turn off cache lines by balancing an estimated power gain against an estimated power loss, we need not also determine whether Takahashi expressly discloses determining whether to power down a cache way.

As detailed above, Petitioner also provides adequate reasons with rational underpinning for why a person of ordinary skill in the art would have been motivated to combine the teachings of Takahashi and Hu to reduce power consumption by efficiently powering down portions of a cache. *See* Pet. 42–46 (citing Ex. 1102 ¶¶ 102–106). There would have been a reasonable expectation of success in doing so, as both references teach using information related to cache misses to determine when to power down a portion of a cache, and Hu additionally teaches turning off cache lines based on a comparison between an estimated power loss due to cache misses and an estimated power gain from turning off a portion of a cache. *See id.* For the reasons discussed above, Patent Owner's arguments do not undermine Petitioner's persuasive showing.

c. Conclusion as to Claim 1

Based on the entire record before us, we determine that Petitioner has shown that the combination of Takahashi and Hu teaches or suggests all the limitations of independent claim 1 and that a person of ordinary skill in the art would have combined the teachings of Takahashi and Hu in the manner asserted with a reasonable expectation of success. "Once all relevant facts are found, the ultimate legal determination [of obviousness] involves the weighing of the fact findings to conclude whether the claimed combination would have been obvious to an ordinary artisan." *Arctic Cat Inc. v. Bombardier Recreational Prods. Inc.*, 876 F.3d 1350, 1361 (Fed. Cir. 2017). On balance, we determine that Petitioner has established by a preponderance of the evidence that claim 1 would have been obvious over the combination of Takahashi and Hu.

2. Claim 12

Independent claim 12 is a device claim with limitations substantially similar to those in method claim 1. *Compare* Ex. 1101, 8:8–18, *with id.* at 7:27–35. For the additional claim language in limitation 12[a], "power switching circuitry adapted to selectively provide power," Petitioner cites Takahashi's power control circuit. 1312 Pet. 44 (citing Ex. 1103, 1:29–4:23, 4:43–46, 6:10–18, 8:27–28, 8:43–45, 9:43–10:31, Figs. 1, 3, 4); *see* Pet. 48 (citing Ex. 1103, 1:58–2:15, 4:43–46, 6:10–18, 8:27–45, 9:37–10:33, Fig. 4).[15] For "power management circuitry adapted to determine whether to power down," as recited in limitation 12[b], Petitioner cites Takahashi's way number control circuit. 1312 Pet. 44 (citing Ex. 1103, 1:29–4:23, 4:39–46, 4:64–5:13, 6:2–10, 8:3–6, 8:12–10:33, Figs. 3, 4); *see* Pet. 49 (citing Ex. 1103, 1:58–2:35, 4:43–46, 6:3–18, 8:51–11:48, Fig. 4). Petitioner has shown sufficiently that Takahashi teaches the claimed "circuitry," which Patent Owner does not specifically dispute. *See* 1312 PO Resp.

The remainder of Petitioner's analysis for claim 12 is substantially the same as its analysis for claim 1. *See* 1312 Pet. 43–54; Pet. 47–50 (referring back to claim 1 analysis). In response, Patent Owner presents substantially the same arguments for claim 12 as it does for claim 1. *See* 1312 PO Resp. 21–34; 1312 PO Sur-reply 15–24.

For the reasons discussed above and for the reasons discussed in our consideration of claim 1, we determine that Petitioner has shown that the combination of Takahashi and Hu teaches or suggests all the limitations of independent claim 12 and that a person of ordinary skill in the art would

[15] Although the Petition in IPR2018-01661 does not challenge claim 12 directly, it includes an analysis of claim 12 because it challenges claims 15 and 16, which depend from claim 12. *See* Pet. 47–50.

have combined the teachings of Takahashi and Hu in the manner asserted with a reasonable expectation of success. On balance, considering the entire record before us, we determine that Petitioner has shown by a preponderance of the evidence that claim 12 would have been obvious over the combination of Takahashi and Hu.

3. Claims 2, 3, 13, and 14

Claims 2 and 13 depend from independent claims 1 and 12, respectively, and require the claimed "component" to be a "cache memory." Ex. 1101, 7:36–37, 8:19–20. Claims 3 and 14 depend from claims 2 and 13, respectively, and further recite that "the component is a cache memory that comprises independently powered portions." *Id.* at 7:38–40, 8:21–23.

Petitioner provides a detailed analysis explaining where both Takahashi and Hu teach these limitations. Pet. 35, 47; 1312 Pet. 37–38, 54–55. For instance, Takahashi "relates to cache memories" and discloses a cache memory comprising a plurality of ways, each of which can be powered down separately. Ex. 1103, 1:9–11, 2:24–38. Hu also discloses turning off individual cache lines in a device. Ex. 1104, 163, 168.

Patent Owner does not specifically dispute that Takahashi and Hu teach the limitations expressly recited by claims 2, 3, 13, and 14. *See* PO Resp.; 1312 PO Resp. Based on the complete record, and for the reasons explained by Petitioner, we find that the combination of Takahashi and Hu teaches all the limitations of claims 2, 3, 13, and 14 and that a person of ordinary skill in the art would have combined the teachings of Takahashi and Hu in the manner asserted with a reasonable expectation of success. On balance, considering the entire record before us, we determine that Petitioner has shown by a preponderance of the evidence that claims 2, 3, 13, and 14 would have been obvious over the combination of Takahashi and Hu.

4. Claims 4, 5, 15, and 16

Claim 4 depends from claim 2 and further recites "wherein the determining is responsive to an amount of information that is stored only in the cache memory." Ex. 1101, 7:41–43. Claim 15 depends from claim 13 and recites a similar limitation. *Id.* at 8:24–26. Claim 5 depends from claim 1 and further recites "wherein the determining is responsive to an amount of information associated with dirty bits." *Id.* at 7:44–46. Claim 16 depends from claim 12 and recites a similar limitation. *Id.* at 8:27–29. Petitioner contends that the combination of Takahashi and Hu teaches these limitations. Pet. 51–56.

As a preliminary matter, Petitioner contends that a person of ordinary skill in the art would have understood that dirty information in a cache is an example of "information that is stored only in the cache memory," as recited in claims 4 and 15, because dirty information is modified information that has not yet been saved to main memory. *Id.* at 51 (citing Ex. 1102 ¶ 122). Petitioner asserts that this understanding is consistent with the description of dirty information in the '014 patent. *See id.* (citing Ex. 1101, 6:37–39 ("Dirty information is information that is stored in the cache but not in another . . . memory unit."")). We agree with Petitioner's position, which Patent Owner does not dispute. *See* PO Resp. 38.

Petitioner also contends that a person of ordinary skill in the art would have understood that cache lines use dirty bits to indicate whether the cache line contains dirty information. Pet. 55 (citing Ex. 1102 ¶ 128; Ex. 1101, 2:2–7; Ex. 1105, 4:24–29); *see also* Ex. 1104, 186 (Hu indicating that "dirty lines" are "distinguished by the dirty bit"). Thus, according to Petitioner, dirty information in a cache is "an amount of information associated with dirty bits," as recited in claims 5 and 16. *See* Pet. 55. We agree with

Petitioner's position, which Patent Owner does not dispute. *See* PO Resp. 38.

Petitioner then argues that Hu teaches determining whether to power down a cache line based on the amount of dirty information in the cache line. Pet. 52–53. In particular, Hu teaches that powering down a cache line containing dirty information consumes more power than powering down a clean line because it leads to additional L1 cache misses, which require additional L2 cache accesses, thereby consuming additional dynamic power. Ex. 1104, 166 ("Turning off a dirty line results in an early writeback . . . ."), 167 ("The *L2Access:leak* ratio relates dynamic energy due to an additional miss (or writeback) to a single clock cycle of static leakage energy in the L1 cache. Multiplying this by the number of extra L2 accesses induced by cache decay gives the dynamic cost induced."), 169 ("[I]f we prematurely turn off a line that may still have hits, then we inject extra misses that incur dynamic power for L2 cache accesses."); *see* Pet. 52. Petitioner contends that Hu teaches considering these additional L2 accesses when determining whether to power down a cache line. Pet. 52–53 (citing Ex. 1104, 167). Therefore, Petitioner contends, Hu teaches determining whether to power down a portion of a component based on the amount of dirty information in the cache lines. *Id.* at 53 (citing Ex. 1102 ¶ 124). Accordingly, Petitioner contends that the combined system of Takahashi and Hu would determine whether to power down a cache line based on Hu's estimation of power loss, which includes consideration of the additional power loss from powering down dirty cache lines. *Id.* (citing Ex. 1102 ¶ 125).

Patent Owner contends that Petitioner's argument "is based on the false allegation that Hu teaches powering down portions of a cache in response to the *L2Access:leak* ratio, where in fact Hu teaches . . . powering

down after elapsing of a fixed decay interval." PO Resp. 38–39. As discussed previously, however, the Federal Circuit found that Hu teaches using the *L2Access:leak* ratio to determine whether to power down a portion of a cache. *See Intel*, 858 Fed. App'x at 355. Thus, we do not agree with Patent Owner's argument.

Patent Owner further argues that the *L2Access* value in the *L2Access:leak* ratio is independent of the amount of dirty information in the cache because it represents the dynamic energy due to an additional cache miss (or writeback), i.e., a single L2 access. PO Resp. 39 (citing Ex. 2107 ¶ 60); PO Sur-reply 23–24. Patent Owner criticizes Petitioner's citation to Hu's statement that "[m]ultiplying [*L2Access*] by the number of extra L2 accesses induced by cache decay gives the dynamic cost induced," arguing that Petitioner has not otherwise alleged that Hu teaches powering down portions of a cache based on the number of extra L2 accesses induced by cache decay. PO Resp. 39 (citing Ex. 1104, 167); *see* Pet. 53.

We agree with Patent Owner that the *L2Access* parameter itself does not reflect the number of L2 accesses that might be caused by turning off a cache line because it represents the dynamic power *per access* to L2 memory. *See* Ex. 1104, 167. Nevertheless, although Hu teaches determining whether to power down a cache line based on the *L2Access:leak* ratio that Hu uses to evaluate cache decay policies, Hu also generally teaches determining whether to turn off a cache line by comparing an estimated power gain (e.g., static L1 cache leakage power saved by turning off a cache line) to an estimated power loss (e.g., extra dynamic power dissipated by accessing the L2 cache). *See* Ex. 1104, 166–69, 173–76. The estimated power loss in this comparison takes into consideration the additional L2

accesses that powering down a cache line containing dirty information would cause. *See id.* at 166–67; Pet. 52–53; Ex. 1102 ¶ 124.

For these reasons, we find that the combination of Takahashi and Hu teaches that determining whether to power down a portion of a cache is responsive to an amount of dirty information, which is "information that is stored only in the cache memory," as recited in claims 4 and 15, and is "information associated with dirty bits," as recited in claims 5 and 16. We also find that a person of ordinary skill in the art would have combined the teachings of Takahashi and Hu in the manner asserted with a reasonable expectation of success. On balance, considering the complete record, we determine that Petitioner has shown by a preponderance of the evidence that claims 4, 5, 15, and 16 would have been obvious over the combination of Takahashi and Hu.

5. Claims 18 and 20

Claim 18, which depends from claim 12, requires the claimed device to be "further adapted to generate cache statistics" and requires the determination whether to power down to be "responsive to the generated cache statistics." Ex. 1101, 8:35–37. Claim 20, which depends from claim 12, requires the claimed device to be "further adapted to monitor integrated circuit behavior" and requires the determination whether to power down to be "responsive to the results of the monitoring." *Id.* at 8:41–43. Petitioner relies on Takahashi for teaching these limitations. 1312 Pet. 55–56 (citing Ex. 1103, 2:36–4:23, 4:35–39, 4:43–46, 5:6–6:2, 6:59–8:11, 9:41–62, Fig. 2); *see id.* at 38–43 (obviousness analysis based on Takahashi alone).

For claim 18, Petitioner cites Takahashi's cache-miss rate measuring circuit, which comprises a cache-miss rate value store queue. *Id.* at 39

(citing Ex. 1103, 6:63–7:2, Fig. 2). The cache-miss rate value store queue calculates cache-miss rates using data from a cache-miss counter and a memory access counter. *See id.* (citing Ex. 1103, 5:64–67, 7:31–37). Petitioner contends that, consistent with the example in the '014 patent of cache-miss rates representing cache statistics used to estimate power loss, the cache-miss rates generated by Takahashi's cache-miss rate value store queue are generated "cache statistics," as recited in claim 18. *Id.* at 39–40 (citing Ex. 1101, 7:3–9; Ex. 1002 ¶ 97).

For claim 20, Petitioner again relies on Takahashi's cache-miss rate measuring circuit. *Id.* at 41–43. Specifically, Petitioner contends that Takahashi's cache-miss rate measuring circuit monitors integrated circuit behavior when it measures the number of cache misses during a memory access operation. *Id.* at 42–43 (citing Ex. 1103, 2:37–50, 3:15–21, 5:5–13; Ex. 1002 ¶ 104).

Based on the evidence of record, we find that Takahashi discloses a circuit that generates cache-miss rates, which are "cache statistics," as recited in claim 18. *See* Ex. 1103, 5:64–67, 6:63–7:2, 7:31–37. We also find that Takahashi's cache-miss rate measuring circuit "monitor[s] integrated circuit behavior," as recited in claim 20, when it measures the number of cache misses. *See id.* at 2:37–50, 3:15–21, 5:5–13. As discussed previously, Takahashi determines the optimal number of cache ways based on changes in the cache-miss rate that would result from powering down a cache way. *Id.* at 2:54–58, 5:67–6:6. Further, Hu teaches comparing an estimated power gain (e.g., static L1 cache leakage power saved by turning off a cache line) with an estimated power loss due to additional cache misses (e.g., extra dynamic power dissipated by accessing the L2 cache) that would result from powering down a cache line and determining when to power

down a cache line by balancing the two. *See* Ex. 1104, 166–69, 173–76. When the teachings of Takahashi and Hu are combined as proposed by Petitioner, we are persuaded that a person of ordinary skill in the art would have used the cache-miss rate, as taught by Takahashi, to determine the estimated power loss based on additional cache misses, as taught by Hu. *See* 1312 Pet. 49–54 (explaining why a person of ordinary skill in the art would have combined Takahashi and Hu with respect to independent claim 12), 56 (analysis of claims 18 and 20 referring to claim 12 motivation to combine analysis); *see also* Pet. 42–46 (explaining why a person of ordinary skill in the art would have combined Takahashi and Hu with respect to independent claim 1).

Patent Owner's only argument with respect to claims 18 and 20 is that Takahashi does not power down a cache way in response to cache-miss rates because Takahashi powers down before measuring the cache-miss rate. 1312 PO Resp. 39 (referring back to analysis of obviousness challenge based on Takahashi alone); *see id.* at 20. As Petitioner points out, Patent Owner raised a similar argument with respect to the independent claims. *See* 1312 Pet. Reply 14, 20. In our discussion of claim 1, we explained that, in view of the ambiguity in Takahashi's disclosure, a person of ordinary skill in the art would not have understood powering down a cache way before measuring a change in the cache-miss rate to be a basic principle of operation that would be altered by combining Takahashi with Hu. *See supra* § II.E.1.b. And as we further explained, we need not determine whether Takahashi expressly discloses determining whether to power down a cache way because Hu teaches determining whether to power down cache lines by balancing an estimated power gain against an estimated power loss. *See id.* Patent Owner's argument that claims 18 and 20 would not have been

obvious over Takahashi and Hu because Takahashi teaches powering down before measuring the cache-miss rate does not undermine Petitioner's persuasive showing for the same reasons discussed previously with respect to the independent claims.

For these reasons, we find that the combination of Takahashi and Hu teaches a device "adapted to generate cache statistics" and "wherein the determination [to power down] is responsive to the generated cache statistics," as recited in claim 18. We also find that the combination of Takahashi and Hu teaches a device "adapted to monitor integrated circuit behavior" and "wherein the determination [to power down] is responsive to the results of the monitoring," as recited in claim 20. We further find that a person of ordinary skill in the art would have combined the teachings of Takahashi and Hu in the manner asserted with a reasonable expectation of success. On balance, we determine that Petitioner has shown by a preponderance of the evidence that claims 18 and 20 would have been obvious over the combination of Takahashi and Hu.

F. Asserted Obviousness over Takahashi, Hu, and Cohen

Petitioner contends that claims 4, 5, 15, and 16 also are unpatentable under 35 U.S.C. § 103(a) for obviousness over the combined teachings of Takahashi, Hu, and Cohen. Pet. 56–66.

As discussed in Section II.E.4, Petitioner contends that a person of ordinary skill in the art would have understood that dirty information in a cache is an example of "information that is stored only in the cache memory," as recited in claims 4 and 15, because dirty information is modified information that has not yet been saved to main memory. Petitioner asserts that Cohen confirms this understanding. Pet. 57 (citing Ex. 1105, 2:61–67 ("[I]f the cache is a writeback cache (data written to

cache is not written to memory until a valid cache line is needed for new data), . . . the cache will contain data not stored elsewhere. This data [is] referred to as 'dirty' data . . . ."); Ex. 1102 ¶ 132). As also discussed in Section II.E.4, Petitioner contends that a person of ordinary skill in the art would have understood that cache lines use dirty bits to indicate whether the cache line contains dirty information and, thus, dirty information in a cache is "an amount of information associated with dirty bits," as recited in claims 5 and 16. In this asserted ground, Petitioner further cites Cohen for teaching that the number of dirty bits in a portion of a cache would be related to the amount of dirty information in the same portion of the cache. *Id.* at 58 (citing Ex. 1105, 4:65–5:4; Ex. 1102 ¶ 136). We agree with Petitioner's contentions, which Patent Owner does not dispute. *See* PO Resp. 41–42.

Petitioner also cites Cohen for providing "specific examples of the extra steps a writeback process requires to save dirty information to main memory before powering off a cache line containing dirty information." Pet. 58 (citing Ex. 1105, 3:65–4:49). Additionally, Petitioner contends that Cohen discloses a "busy bit counter" that counts the number of dirty bits in a portion of a cache. *Id.* (citing Ex. 1105, 4:65–5:4). Petitioner contends that a person of ordinary skill in the art "reading Hu in light of Cohen would have understood that the more dirty information there is" (e.g., as indicated by Cohen's busy bit counter), "the more additional writebacks occur, and therefore the more energy is expended to power down the corresponding cache lines." *Id.* at 58–59 (emphasis omitted) (citing Ex. 1104, 167; Ex. 1105, 4:18–23, 4:51–52; Ex. 1102 ¶ 137); Pet. Reply 21. Therefore, Petitioner contends, Hu combined with Cohen teaches that the determination

whether to power down a portion of a component is based on the amount of dirty information in the cache. Pet. 59 (citing Ex. 1102 ¶ 138).

Petitioner also argues that a person of ordinary skill in the art would have been motivated to combine Cohen with Takahashi and Hu. *Id.* at 60–63. Petitioner contends that Cohen, like Takahashi and Hu, is directed to the same problem of reducing power consumption in a cache memory. *Id.* at 60 (citing Ex. 1105, code (57), 1:29–31; Ex. 1103, 2:23–35; Ex. 1104, 161 (Abstract); Ex. 1102 ¶ 140). Petitioner also contends that Cohen similarly teaches solving the problem by powering down portions of a cache. *Id.* at 61 (citing Ex. 1105, 1:36–40, 2:11–13; Ex. 1102 ¶ 141). Petitioner contends that a person of ordinary skill in the art would have been motivated specifically to combine Cohen with Takahashi and Hu because Cohen teaches determining whether to power down a portion of a cache by analyzing cache hits or misses. *Id.* at 61–62 (citing Ex. 1105, 5:61–66, 6:21–30; Ex. 1102 ¶ 142). Moreover, Petitioner contends, a person of ordinary skill in the art would have read Hu in light of Cohen's disclosures regarding specific writebacks involved in powering down a cache that contains dirty information. *Id.* at 62 (citing Ex. 1104, 166, 167, 186; Ex. 1105, 3:65–4:49; Ex. 1102 ¶ 143). Petitioner contends a person of ordinary skill in the art would have had a reasonable expectation of success in combining Takahashi, Hu, and Cohen due to their similarities. *Id.* at 63 (citing Ex. 1102 ¶ 145).

In response, Patent Owner presents the same arguments regarding Hu that it made in connection with the obviousness challenge to claims 4, 5, 15, and 16 based on Takahashi and Hu. *See* PO Resp. 40–41. We disagree with these arguments for the reasons discussed in Section II.E.4. We also disagree with Patent Owner's argument that Cohen fails to remedy the

deficiency in Hu because we do not find Hu to be deficient. *See* PO Resp. 42. Patent Owner offers no argument regarding the motivation to combine Cohen with Takahashi and Hu separate from its argument regarding the combination of Takahashi and Hu. *See id.* at 40–42; Pet. Reply 22–23 (citing Ex. 1112, 81:12–25 (Dr. Wolfe testifying that his opinion regarding motivation to combine Takahashi, Hu, and Cohen is the same as his opinion regarding motivation to combine Takahashi and Hu)).

Based on the evidence of record, and for the reasons explained in the Petition, we find that the combination of Takahashi, Hu, and Cohen teaches the limitations of claims 4, 5, 15, and 16. *See* Pet. 56–59. Petitioner also provides adequate reasons with rational underpinning for why a person of ordinary skill in the art would have been motivated to combine the teachings of Takahashi, Hu, and Cohen with a reasonable expectation of success. *See id.* at 60–63. On balance, we determine that Petitioner has shown by a preponderance of the evidence that claims 4, 5, 15, and 16 would have been obvious over the combined teachings of Takahashi, Hu, and Cohen.

G. Asserted Obviousness over Takahashi, Hu, and Gunther

Petitioner contends that claim 20 also is unpatentable under 35 U.S.C. § 103(a) for obviousness over the combined teachings of Takahashi, Hu, and Gunther. 1312 Pet. 67–69.

In light of our determination that claim 20 is unpatentable over Takahashi and Hu, we decline to address this additional ground. *See SAS Inst. v. Iancu*, 138 S. Ct. 1348, 1359 (2018) (holding that a petitioner "is entitled to a final written decision addressing all of the claims it has challenged"); *Boston Sci. Scimed, Inc. v. Cook Grp. Inc.*, 809 F. App'x 984, 990 (Fed. Cir. 2020) (nonprecedential) (stating that the "Board need not address issues that are not necessary to the resolution of the proceeding,"

such as "alternative arguments with respect to claims [the Board] found unpatentable on other grounds"); *see also Beloit Corp. v. Valmet Oy*, 742 F.2d 1421, 1423 (Fed. Cir. 1984) (An administrative agency is at liberty to reach a decision based on a dispositive issue because doing so "can not only save the parties, the [agency], and [the reviewing] court unnecessary cost and effort," but "can greatly ease the burden on [an agency] commonly faced with a . . . proceeding involving numerous complex issues and required by statute to reach its conclusion within rigid time limits."); *SK Hynix Inc. v. Netlist, Inc.*, IPR2017-00692, Paper 25 at 40 (PTAB July 5, 2018) (determining all challenged claims to be unpatentable and not addressing additional grounds).

H. Credibility of Dr. Raghunathan

Patent Owner asserts that Dr. Raghunathan's testimony should be discounted due to his "undisclosed bias, misunderstanding of legal standards, and clear instances of concealment and misdirection." PO Resp. 2; 1312 PO Resp. 2. To support its allegations, Patent Owner directs our attention to certain facts brought out during Patent Owner's cross-examination of Dr. Raghunathan that purportedly show Dr. Raghunathan's bias resulting from his numerous financial, professional, and personal connections to Petitioner. *See* PO Resp. 43–53; 1312 PO Resp. 43–53.

To assess the probative value of an expert opinion, we must consider the following factors: (1) the interest of the expert in the outcome of the case, (2) the presence or absence of factual evidence supporting the expert's opinion, and (3) the strength of any opposing evidence. *See Ashland Oil, Inc. v. Delta Resins & Refractories, Inc.*, 776 F.2d 281, 294–95 (Fed. Cir. 1985). An additional factor for consideration is "the nature of the matter

sought to be established" (*see id.* at 294), which, in this context, is obviousness.

We agree that Dr. Raghunathan's financial interests and relationships with Petitioner are factors in deciding how much weight to give his testimony. *See, e.g.*, PO Resp. 44–49. Given the extensive and lengthy financial support Intel provides Dr. Raghunathan, we cannot say that Dr. Raghunathan is an entirely disinterested party to the proceeding. *See, e.g.*, Ex. 2108, 29:9–31:22, 32:21–35:20, 36:10–39:6 (identifying hundreds of thousands of dollars of funding provided from Intel to Dr. Raghunathan and other faculty at Purdue University), 40:20–21 ("I have regularly received funds from Intel"), 44:18–20 (stating "Intel has been . . . funding my research for a long time and . . . might fund my research again"), 41:24–42:21 (discussing purchase of Intel stock), 75:5–76:20 (discussing research projects involving Dr. Raghunathan and Intel employees), 84:17–88:16 (discussing assignment of patent rights from Dr. Raghunathan to Intel).

Although "the opinion testimony of a party having a direct interest in the pending [proceeding] is less persuasive than opinion testimony by a disinterested party, it cannot be disregarded for that reason alone and may be relied upon when sufficiently [persuasive]." *Ashland Oil*, 776 F.2d at 294. In this context, we have considered the persuasiveness of Dr. Raghunathan's testimony, including the specific testimony Patent Owner contends is evidence of his bias.

For example, Patent Owner contends that Dr. Raghunathan misunderstands or disregards the appropriate legal standard for the level of skill in the art, asserting that Dr. Raghunathan takes the position that a level of skill requiring a B.S. degree might actually require an equivalent or even *higher* level of skill than one requiring an M.S. degree. *See* PO Resp. 49

(citing Ex. 2108, 103:19, 105:14–21, 106:11–19); 1312 PO Resp. 49 (citing same). We disagree with Patent Owner.

The cited testimony relates to Patent Owner's questioning as to whether Dr. Raghunathan's proposed level of ordinary skill in the art is higher or lower than the level of skill adopted by the Board. In his declarations, Dr. Raghunathan asserted that, due to the complexity in designing power management systems, a person of ordinary skill in the art would have had at least an M.S. degree in certain fields and three years of experience with computer hardware and software power management design techniques. Ex. 1102 ¶¶ 57–58; Ex. 1002 ¶¶ 59–60. Dr. Raghunathan explained that a person with an M.S. degree in the relevant field would have been involved with "research and/or experiments involving computer hardware and software power management design techniques." Ex. 1102 ¶ 58; Ex. 1002 ¶ 60. Dr. Raghunathan further explained that a person having at least three years of "academic or industry experience in computer or industry experience in computer hardware and software power management design techniques" would have had additional experience with designing real-world computer hardware and software-based power management systems. Ex. 1102 ¶ 58; Ex. 1002 ¶ 60. In our Decisions on Institution, we adopted Dr. Raghunathan's proposal with a modification as to the level of education, determining that a person of ordinary skill in the art would have had "at least a B.S. degree in computer science [or other relevant field,] and three years of experience with computer hardware and software power management design techniques." 1661 Inst. Dec. 7; 1312 Inst. Dec. 7. We stated that we were "not persuaded a person of ordinary skill in the art would have needed an M.S. degree in addition to a B.S. degree if such a person also

had three years of relevant industry experience (e.g., experience with power management techniques)." 1661 Inst. Dec. 7; 1312 Inst. Dec. 7.

When questioned as to whether the level of skill adopted by the Board or the level of skill proposed in his declarations requires more skill, Dr. Raghunathan testified that he could not give an answer with certainty but explained that getting a master's degree typically requires more years of education than getting a bachelor's degree. Ex. 2108, 103:14–23, 104:20–25. Dr. Raghunathan also declined to state that any specific individual who has a master's degree is more skilled than a specific individual who has a bachelor's degree. Ex. 2108, 101:19–102:9, 104:2–105:8. Thus, contrary to Patent Owner's assertions, Dr. Raghunathan did not state that a level of skill requiring a B.S. degree might actually require an equivalent or even *higher* level of skill than one requiring an M.S. degree. We also disagree with Patent Owner's contentions that Dr. Raghunathan misunderstood or disregarded the appropriate legal standard for the level of skill in the art. We have reviewed Dr. Raghunathan's testimony and determine that Dr. Raghunathan has demonstrated sufficient understanding of the legal standards relating to the level of skill in the art that he applied when rendering his testimony. *See, e.g.*, Ex. 2108, 98:5–106:19; Ex. 1002 ¶¶ 55–60, 62; Ex. 1102 ¶¶ 53–58, 60.

Patent Owner also contends that Dr. Raghunathan's bias is reflected in lack of clarity and his alleged mischaracterizations of Takahashi and Hu. PO Resp. 49–50 (citing Ex. 2108, 132:25–133:12, 151:10–152:5), 52 (citing Ex. 2108, 137:10–139:5; Ex. 1102 ¶¶ 71–80), 52–53 (citing Ex. 2108, 168:15–169:10, 171:8–18; Ex. 1102 ¶ 100). On this point, we note that Patent Owner was afforded the opportunity to provide contrary expert testimony and to test Dr. Raghunathan's testimony through

cross-examination, both of which are traditional mechanisms to guard against deceptive testimony and which we are in a position to evaluate. The persuasiveness of Dr. Raghunathan's testimony as to the teachings of the asserted art is considered in our obviousness analysis. *See infra* §§ II.E, II.F; *see also* 1661 Final Dec. §§ II.D.3, II.E.3; 1312 Final Dec. §§ II.D.3, II.E.3.

We have also considered Patent Owner's allegations of "concealment and misdirection" in determining how much weight to accord Dr. Raghunathan's testimony. *See* PO Resp. 43, 50–53; 1312 PO Resp. 43, 50–53. We note that some of the allegedly concealed information was, in fact, disclosed on Dr. Raghunathan's curriculum vitae (CV).[16] For example, Patent Owner asserts "material information left undisclosed" by Dr. Raghunathan includes "[a]t least one former graduate student who subsequently obtained full-time employment with the Petitioner" and "[a]dditional former students who were employed as interns by Petitioner." PO Resp. 51 (citing Ex. 2108, 51:17–21, 56:5–57:23); 1312 PO Resp. 51 (citing same). Dr. Raghunathan's declarations, however, state that "former student advisees have joined teams at Intel." Ex. 1002 ¶ 4; Ex. 1102 ¶ 4.

Patent Owner further alleges that Dr. Raghunathan failed to identify patent application WO 2018/034681, which purportedly is owned by Petitioner and lists Dr. Raghunathan as well as thirteen of Petitioner's employees as coinventors. PO Resp. 52 (citing Ex. 2108, 86:10–22; Ex. 2112); 1312 PO Resp. 52 (citing same). Dr. Raghunathan's CV, however, explicitly states that "[o]nly US patents [are] listed below."

[16] As noted by Patent Owner, Petitioner did not attach a copy of Dr. Raghunathan's CV to his declaration submitted in IPR2018-01661 but did attach a copy to his declaration submitted in IPR2018-01312. PO Resp. 45 n.8.

Ex. 1002, 88. We decline to infer that Dr. Raghunathan's failure to include patent application WO 2018/034681 in a list of issued US patents constitutes concealment or misdirection.

Patent Owner further alleges that Dr. Raghunathan failed to disclose "gifts" and other payments by Intel to fund Dr. Raghunathan's research. PO Resp. 51 (citing Ex. 2108, 30:8–15, 36:10–37:11, 38:7–39:17); 1312 PO Resp. 51 (citing same). Some of this funding, which appears to have been provided by Intel to Purdue University and not directly to Dr. Raghunathan, is listed on Dr. Raghunathan's CV. *See* Ex. 2108, 30:8–15; Ex. 1002, 86. Other funding appears to have been provided after Dr. Raghunathan's declarations were signed. *See* Ex. 1002, 85–88; Ex. 2108, 36:10–37:11, 38:7–39:17. As acknowledged by Patent Owner, Dr. Raghunathan's CV does list numerous instances of funding provided directly by Intel to Dr. Raghunathan. *See* PO Resp. 44; Ex. 1002, 85–88.

We have also reviewed the remaining testimony that Patent Owner contends is evidence of Dr. Raghunathan's misdirection and concealment, including Dr. Raghunathan's alleged failure to disclose employment or other funding or gifts from Intel or lack of recollection as to certain matters. *See, e.g.*, PO Resp. 51 (citing Ex. 2108, 28:10–12, 111:4–11, 114:20–115:2), 53 (citing Ex. 2108, 33:7–34:2, 69:9–71:3, 118:5–119:15, 177:7–10); 1312 PO Resp. 51, 53 (citing same). We disagree with Patent Owner's contentions that such testimony is evidence of willful concealment or misdirection.

For example, Patent Owner cites to Dr. Raghunathan's deposition testimony at 111:4–11 and 114:20–115:2 as evidence that Dr. Raghunathan failed to disclose gifts totaling over $90,000. *See* PO Resp. 51; 1312 PO Resp. 51. However, the testimony at these pages relates to conversations Dr. Raghunathan had with attorneys, not to any funding received from

Petitioner. Ex. 2108, 111:4–11, 114:20–115:2. Patent Owner also cites to Dr. Raghunathan's deposition testimony at 28:10–12 as evidence that Dr. Raghunathan failed to disclose "his employment by Petitioner for a full year within the last five years." PO Resp. 51; 1312 PO Resp. 51. However, Dr. Raghunathan did not testify that he was a full-time employee of Petitioner. Rather, Dr. Raghunathan stated he was a technical consultant to Petitioner for approximately twelve months spanning 2015 and 2016. Ex. 2108, 28:10–12. Dr. Raghunathan further explained that this consulting work was part time as his employer, Purdue University, limited the amount of outside consultant work that its faculty may perform to no more than 40 days a year and that his consulting work fell well within those guidelines. *See id.* at 43:17– 44:9. Given that Dr. Raghunathan's curriculum vitae does not detail part time consulting work, and Dr. Raghunathan disclosed his consulting work when asked to describe his relationship with Intel, we do not find the cited testimony to constitute evidence of concealment of misdirection.

Finally, Patent Owner cites to Dr. Raghunathan's alleged "lack of recollection as to pertinent facts in this case" as evidence of lack of candor. PO Resp. 53; 1312 PO Resp. 53 (citing Ex. 2108, 33:7–34:2, 69:9–71:3, 118:5–119:15, 177:7–10). However, the cited testimony appears to show Dr. Raghunathan provided his recollection of facts known to him at the time of his deposition as opposed to demonstrating a lack of candor.

For the reasons explained above, we find insufficient basis to disregard Dr. Raghunathan's opinions based solely on Patent Owner's allegations of bias, concealment, and misdirection; rather, we accord Dr. Raghunathan's opinions due weight in reaching the determinations made in this Decision.

I. Constitutional Challenges

Patent Owner argues that "retroactive application" of *inter partes* review to the '014 patent "would violate the Due Process Clause of the Fifth Amendment, because the '014 patent issued and claims inventions that were publicly disclosed under the Patent Act well before the passage or effective date of the Leahy-Smith America Invents Act of 2011." PO Resp. 54. Patent Owner also argues that "[b]ecause the panel has not yet been appointed by the President and confirmed by the Senate, the challenged claims should not be invalidated." *Id.* at 56. We decline to consider Patent Owner's constitutional challenges, as the issues have been addressed in *Celgene Corp. v. Peter*, 931 F.3d 1342 (Fed. Cir. 2019), and *United States v. Arthrex, Inc.*, 141 S. Ct. 1970 (2021).

III. CONCLUSION[17]

For the foregoing reasons, we determine Petitioner has shown by a preponderance of the evidence that claims 1–5, 12–16, 18, and 20 are unpatentable, summarized in the following tables:

IPR2018-01661:

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not Shown Unpatentable
1–5, 15, 16	103(a)	Takahashi, Hu	1–5, 15, 16	
4, 5, 15, 16	103(a)	Takahashi, Hu, Cohen	4, 5, 15, 16	
Overall Outcome			1–5, 15, 16	

IPR2018-01312:

Claim(s)	35 U.S.C. §	Reference(s)/ Basis	Claims Shown Unpatentable	Claims Not Shown Unpatentable
12–14, 18, 20	103(a)	Takahashi, Hu	12–14, 18, 20	
20	103(a)	Takahashi, Hu, Gunther[18]		
Overall Outcome			12–14, 18, 20	

[17] Should Patent Owner wish to pursue amendment of the challenged claims in a reissue or reexamination proceeding subsequent to the issuance of this decision, we draw Patent Owner's attention to the April 2019 *Notice Regarding Options for Amendments by Patent Owner Through Reissue or Reexamination During a Pending AIA Trial Proceeding*. *See* 84 Fed. Reg. 16654 (Apr. 22, 2019). If Patent Owner chooses to file a reissue application or a request for reexamination of the challenged patent, we remind Patent Owner of its continuing obligation to notify the Board of any such related matters in updated mandatory notices. *See* 37 C.F.R. §§ 42.8(a)(3), (b)(2).

[18] As explained above, in light of our determination that claim 20 is unpatentable for obviousness over Takahashi and Hu, we decline to address this additional ground.

IV. ORDER

Accordingly, it is

ORDERED that the Final Written Decisions in IPR2018-01661 and IPR2018-01312 are modified to include this Remand Decision but are not otherwise modified on remand;

FURTHER ORDERED that claims 1–5, 12–16, 18, and 20 of U.S. Patent No. 8,020,014 B2 have been shown to be unpatentable; and

FURTHER ORDERED that, because this is a final written decision, parties to this proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

FOR PETITIONER:

Donald R. Steinberg
Dominic E. Massa
Daniel V. Williams
WILMER CUTLER PICKERING HALE AND DORR LLP
Don.Steinberg@wilmerhale.com
Dominic.Massa@wilmerhale.com
Daniel.Williams@wilmerhale.com

FOR PATENT OWNER:

Michael R. Fleming
Benjamin Hattenbach
Amy E. Proctor
Kamran Vakili
IRELL & MANELLA LLP
mfleming@irell.com
bhattenbach@irell.com
aproctor@irell.com
kvakili@irell.com

US008020014B2

(12) **United States Patent**
Priel et al.

(10) **Patent No.: US 8,020,014 B2**
(45) **Date of Patent: Sep. 13, 2011**

(54) **METHOD FOR POWER REDUCTION AND A DEVICE HAVING POWER REDUCTION CAPABILITIES**

(75) Inventors: **Michael Priel**, Hertzelia (IL); **Dan Kuzmin**, Givat Shmuel (IL); **Anton Rozen**, Gedera (IL); **Leonid Smolyanski**, Or Akiva (IL)

(73) Assignee: **Freescale Semiconductor, Inc.**, Austin, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 867 days.

(21) Appl. No.: **11/914,079**

(22) PCT Filed: **May 11, 2005**

(86) PCT No.: **PCT/IB2005/051539**

§ 371 (c)(1),
(2), (4) Date: **Nov. 9, 2007**

(87) PCT Pub. No.: **WO2006/120507**

PCT Pub. Date: **Nov. 16, 2006**

(65) **Prior Publication Data**

US 2008/0209248 A1 Aug. 28, 2008

(51) **Int. Cl.**
G06F 1/32 (2006.01)
G06F 1/00 (2006.01)
G06F 11/30 (2006.01)

(52) **U.S. Cl.** **713/320**; 713/300; 713/340

(58) **Field of Classification Search** 713/340
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,483,656 A * 1/1996 Oprescu et al. 713/320
6,169,419 B1 1/2001 De et al.
6,477,388 B1 * 11/2002 Schmutz 455/522
6,742,100 B1 * 5/2004 Schnee et al. 711/173
7,203,852 B2 * 4/2007 Cohen et al. 713/320
2001/0038277 A1 11/2001 Burstein et al.
2002/0000797 A1 1/2002 Schultz et al.
2002/0049918 A1 4/2002 Kaxiras et al.
2003/0139927 A1 7/2003 Gabarra et al.
2003/0145241 A1 7/2003 Hu et al.
2004/0008056 A1 1/2004 Kursun et al.
2004/0025068 A1 2/2004 Gary et al.
2004/0052098 A1 3/2004 Burtsein et al.
2004/0133746 A1 7/2004 Edirisooriya et al.
2004/0158750 A1 8/2004 Syed et al.
2004/0230849 A1 11/2004 Dhong et al.
(Continued)

OTHER PUBLICATIONS

Burd et al; "A Dynamic Voltage Scaled Microprocessor System"; IEEE Journal of Solid-State Circuits, vol. 35, No. 11, Nov. 2000.
(Continued)

Primary Examiner — Suresh K Suryawanshi

(57) **ABSTRACT**

A method for power reduction, the method includes determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode, and selectively providing power to at least a portion of a component of an integrated circuit during a low power mode. A device having power reduction capabilities, the device includes power switching circuitry, and a power management circuitry adapted to determine whether to power down at least the portion of the component during a low power mode in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power.

23 Claims, 2 Drawing Sheets

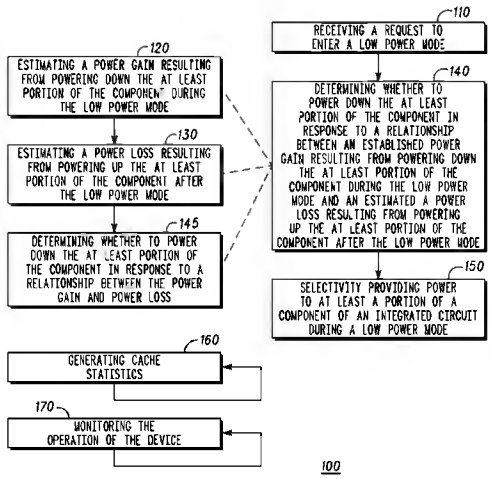

U.S. PATENT DOCUMENTS

2004/0260957 A1 12/2004 Jeddeloh et al.
2005/0044448 A1 2/2005 Verdum

OTHER PUBLICATIONS

Kalla, P. et al., "LRU-SEQ: A Novel Replacement Policy for Transition Energy Reduction in Instruction Caches," IEEE/ACM International Conference on Computer Aided Design; Nov. 9, 2003; ISBN: 1-58113-762-1; pp. 518-522.

Lorch, J. et al., "Software Strategies for Portable Computer Energy Management," IEEE Personal Communications, IEEE Communications Society, US, vol. 5, No. 3; Jun. 1998; pp. 60-73.

PCT International Search Report mailed Apr. 4, 2006 for PCT/IB2005/051539, 4 pages.

* cited by examiner

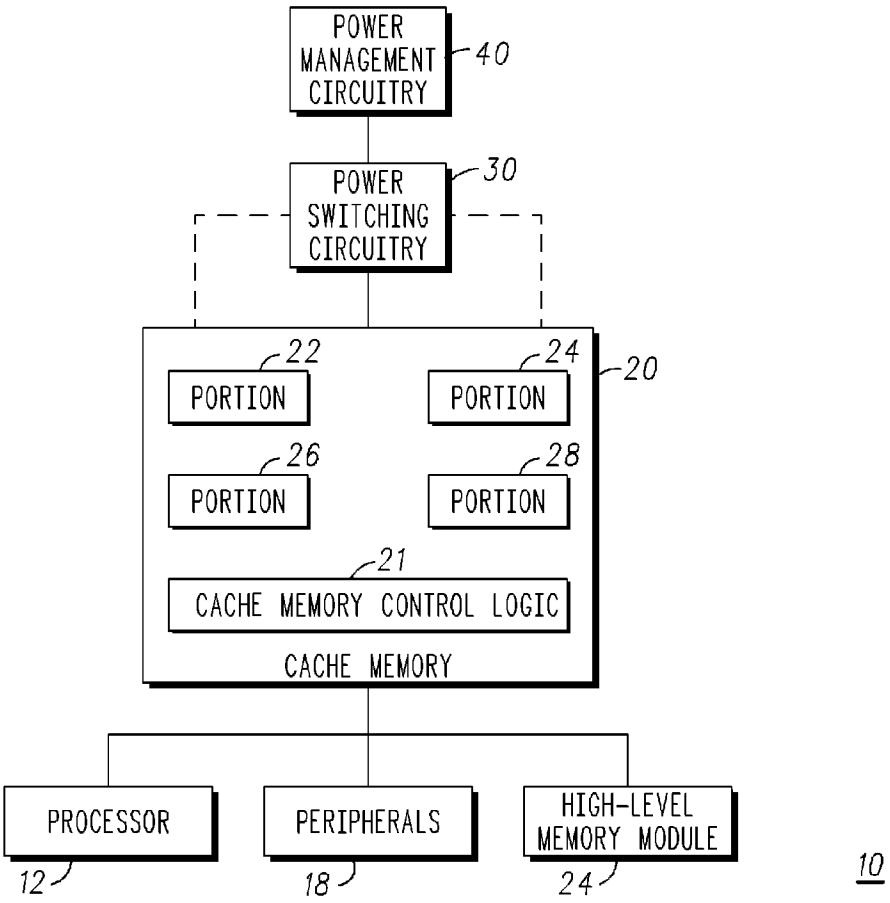
POWER MANAGEMENT CIRCUITRY
40
POWER SWITCHING CIRCUITRY
30
22
PORTION
24
PORTION
20
26
PORTION
28
PORTION
21
CACHE MEMORY CONTROL LOGIC
CACHE MEMORY
PROCESSOR
12
PERIPHERALS
18
HIGH-LEVEL MEMORY MODULE
24
10

FIG. 1

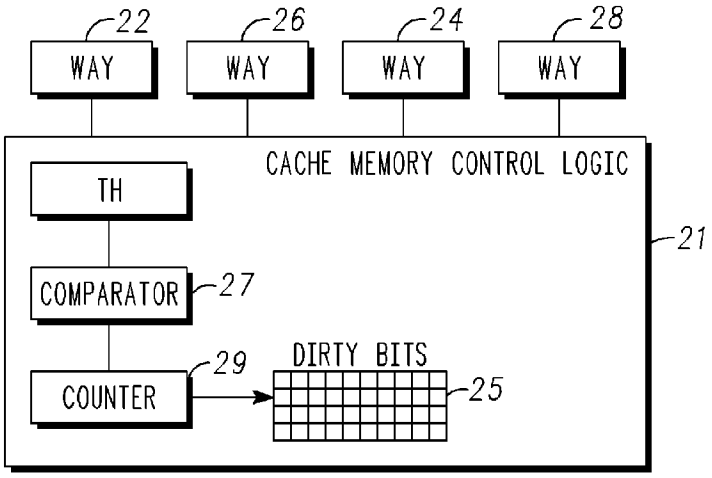
22
WAY
26
WAY
24
WAY
28
WAY
CACHE MEMORY CONTROL LOGIC
TH
COMPARATOR
27
COUNTER
29
DIRTY BITS
25
21
20

FIG. 2

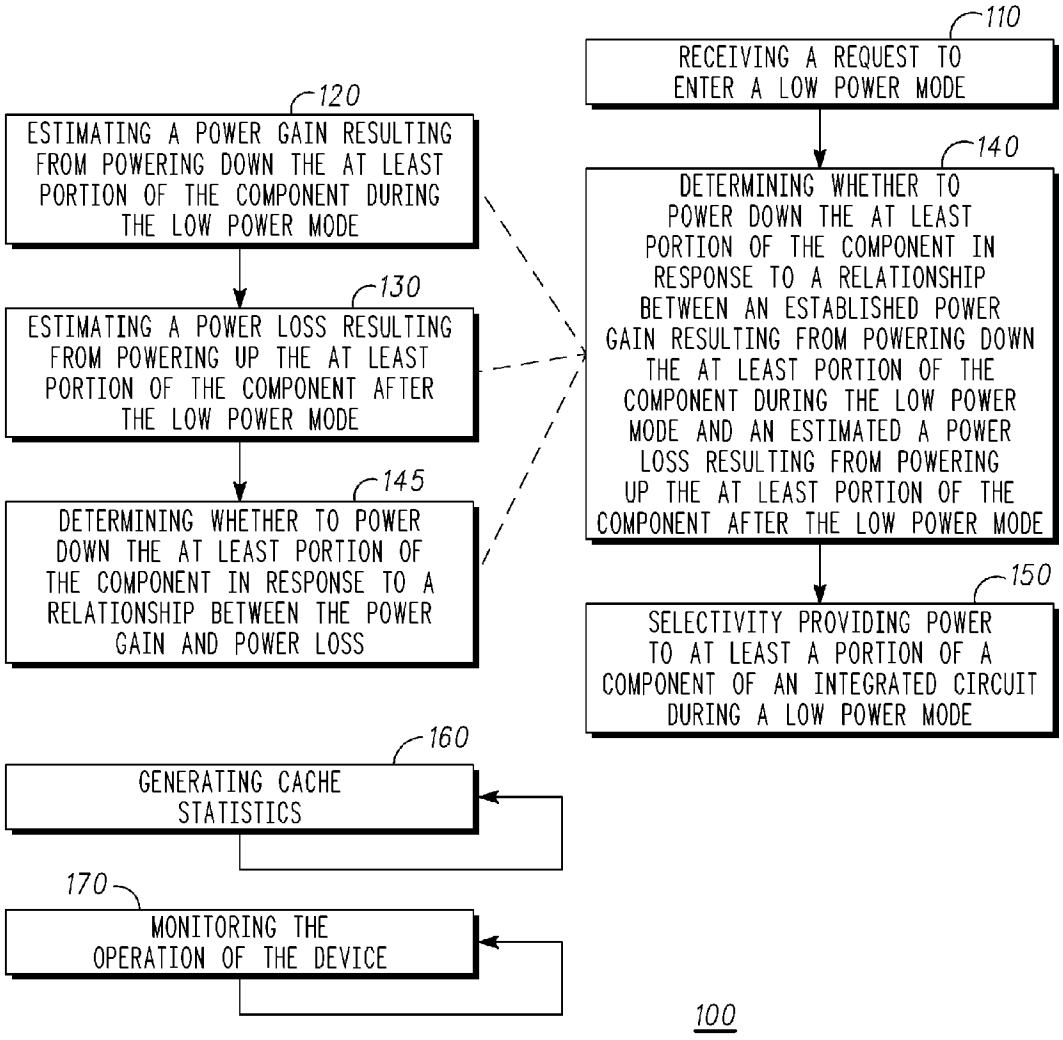
110
RECEIVING A REQUEST TO ENTER A LOW POWER MODE
120
ESTIMATING A POWER GAIN RESULTING FROM POWERING DOWN THE AT LEAST PORTION OF THE COMPONENT DURING THE LOW POWER MODE
130
ESTIMATING A POWER LOSS RESULTING FROM POWERING UP THE AT LEAST PORTION OF THE COMPONENT AFTER THE LOW POWER MODE
145
DETERMINING WHETHER TO POWER DOWN THE AT LEAST PORTION OF THE COMPONENT IN RESPONSE TO A RELATIONSHIP BETWEEN THE POWER GAIN AND POWER LOSS
140
DETERMINING WHETHER TO POWER DOWN THE AT LEAST PORTION OF THE COMPONENT IN RESPONSE TO A RELATIONSHIP BETWEEN AN ESTABLISHED POWER GAIN RESULTING FROM POWERING DOWN THE AT LEAST PORTION OF THE COMPONENT DURING THE LOW POWER MODE AND AN ESTIMATED A POWER LOSS RESULTING FROM POWERING UP THE AT LEAST PORTION OF THE COMPONENT AFTER THE LOW POWER MODE
150
SELECTIVITY PROVIDING POWER TO AT LEAST A PORTION OF A COMPONENT OF AN INTEGRATED CIRCUIT DURING A LOW POWER MODE
160
GENERATING CACHE STATISTICS
170
MONITORING THE OPERATION OF THE DEVICE
100

FIG. 3

METHOD FOR POWER REDUCTION AND A DEVICE HAVING POWER REDUCTION CAPABILITIES

FIELD OF THE INVENTION

The present invention relates to a method for power management and a device having power management capabilities, and especially for power reduction of a cache memory during a low power mode.

BACKGROUND OF THE INVENTION

Cache Memories

Cache modules are high-speed memories that facilitate fast retrieval of information including data and instructions. Typically, cache modules are relatively expensive and are characterized by a small size, especially in comparison to higher-level memory modules.

The performance of modern processor-based systems usually depends upon the cache module performances and especially to a relationship between cache hits and cache misses. A cache hit occurs when an information unit that is present in a cache module memory is requested. A cache miss occurs when the requested information unit is not present in the cache module and has to be fetched from a another memory that is termed a higher-level memory module.

Various cache module modules and processor architectures, as well as data retrieval schemes, were developed over the years, to meet increasing performance demands. These architectures included multi-port cache modules, multi-level cache module architecture, super scalar type processors and the like.

The following U.S patents and patent applications, all being incorporated herein by reference, provide a brief summary of some state of the art cache modules and data fetch methods: U.S. Pat. No. 4,853,846 of Johnson et al., titled "Bus expander with logic for virtualizing single cache control into dual channels with separate directories and prefetch for different processors"; U.S. patent application publication number 20020069326 of Richardson et al., titled "Pipelines non-blocking level two cache system with inherent transaction collision-avoidance"; U.S. Pat. No. 5,742,790 of Kawasaki titled "Detection circuit for identical and simultaneous access in a parallel processor system with a multi-way multi-port cache"; U.S. Pat. No. 6,081,873 of Hetherington et al., titled "In-line bank conflict detection and resolution in a multi-ported non-blocking cache"; U.S and U.S. Pat. No. 6,272,597 of Fu et al., titled "Dual-ported, pipelined, two level cache system".

Processors and other information requesting components are capable of requesting information from a cache module and, alternatively or additionally, from another memory module that can be a higher-level memory module. The higher-level memory module can also be a cache memory, another internal memory and even an external memory.

There are various manners to write information into a cache module or a higher-level memory module. Write-through involves writing one or more information units to the cache module and to the higher-level memory module substantially simultaneously. Write-back involves writing one or more information units to the cache module. The cache module sends one or more updated information units to the high-level memory once that one or more updated information units are removed from the cache. The latter operation is also known in the art as flushing the cache.

Some prior art cache modules include multiple lines that in turn are partitioned to segments. Each segment is associated with a validity bit and a dirty bit. A valid bit indicates whether a certain segment includes valid information. The dirty bit indicates if the segment includes valid information that was previously updated but not sent to the higher-level memory module. If a write back policy is implemented only the segments that are associated with an asserted dirty bit are written to the high-level memory module.

Power Reduction Techniques

Mobile devices, such as but not limited to personal data appliances, cellular phones, radios, pagers, lap top computers, and the like are required to operate for relatively long periods before being recharged. These mobile devices usually includes one or more processors as well as multiple memory modules and other peripheral devices.

In order to reduce the power consumption of mobile devices various power consumption control techniques were suggested. A first technique include selectively shutting down components. Another technique involves reducing the clock frequency of the mobile device. A third technique is known as dynamic voltage scaling (DVS) or alternatively is known as dynamic voltage and frequency scaling (DVFS) and includes altering the voltage that is supplied to a processor as well as altering the frequency of a clock signal that is provided to the processor in response to the computational load demands (also referred to as throughput) of the processor. Higher voltage levels are associated with higher operating frequencies and higher computational load but are also associated with higher energy consumption.

The following U.S. patents, U.S. patent applications, and article, all being incorporated herein by reference, illustrate various power reduction techniques: U.S. patent application publication number 2004/0008056 of Kursun et al.; U.S. Pat. No. 6,169,419 of De et al., U.S. patent application publication number 20040052098 of Burstein et al.; U.S. patent application publication number 20030139927 of Gabara, et al.; U.S. patent application publication number 20040260957 of Jeddeloh et al.; U.S. patent application publication number 20020000797 of Burstein et al.; U.S. patent application publication number 20040025068 of Gary et al.; U.S. patent application publication number 20040158750 of Syed et al.; U.S. patent application publication number 20040230849 of Dhong et al.; U.S. patent application publication number 20010038277 of Burstein et al.; U.S. patent application publication 20050044448 of Verdun; and "A Dynamic Voltage Scaled Microprocessor System", T. D. Burd, T. A. Pering, A. J. Stratakos and R. W. Brodersen, IEEE Journal Journal of solid-state circuits, Vol. 35, No. 11, November 200, all being incorporated herein by reference, provide a brief review of some dynamic voltage scaling techniques.

U.S. patent application publication number 20040133746 of Edirisooriya et al, which is incorporated herein by reference, describes a cache that includes multiple ways. Dirty data can be written only to some predefined ways. These predefined ways are not powered down during low power modes while the other ways of the cache memory are powered down.

U.S. patent application publication number 20030145241 of Hu et al., and U.S. patent application publication number 20020049918 of Kaxiras et al., describe complex methods and apparatuses for powering down cache lines that were not accessed during a cache line decay interval. Thus, cache lines that were not accessed for a long time will be power down during a low power mode.

There is a need to provide an efficient method and device for power reduction.

SUMMARY OF THE PRESENT INVENTION

A method for power reduction and a device having power reduction capabilities, as described in the accompanying claims.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description taken in conjunction with the drawings in which:

FIG. **1** illustrates a device according to an embodiment of the invention;

FIG. **2** illustrates a cache and a power management circuitry, according to an embodiment of the invention; and

FIG. **3** is a flow chart of a power management method, according to an embodiment of the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

The following figures illustrate exemplary embodiments of the invention. They are not intended to limit the scope of the invention but rather assist in understanding some of the embodiments of the invention. It is further noted that all the figures are out of scale.

For convenience of explanation the following detailed explanation relates to a cache memory but the invention can be applied to other components, and especially to components that are characterized by a varying power loss and/or a varying power gain associated with entering a low power mode.

Power loss resulting from entering a low power mode can be associated with operations to be executed before entering the low power mode and/or can be associated with operations that are executed when or after exiting the low power mode.

The power losses or power gain can change over time in a significant manner. The change can be attributed to a change in the number or operations (such as data retrieval operations) associated with entering an/or exiting a low power mode.

If the component to be powered down during a low power mode is a cache or a cache portion then the power loss or gain is conveniently responsive to the power consumed during a flushing operation as well as data retrieval after the cache exits the low power mode. The change can also be attributed to a change in characteristics (such as length) of power down periods.

According to another embodiment of the invention the powering down can be applied in addition to other power saving stages such as clock frequency reduction, dynamic body biasing and the like.

It is assumed, for convenience of explanation that when a high level memory is powered during the low power mode, but this is not necessarily so. It can be checked to determine whether it includes relevant information and if the answer is negative it can also be powered down.

FIG. **1** illustrates a device **10** according to an embodiment of the invention. It is noted that the invention can be applied in devices of different configurations.

Device **10** includes multiple components, that are conveniently located on the same integrated circuit. Device **10** includes a processor **12**, a cache memory **20**, a high level memory module **24**, power switching circuitry **30**, power management circuit **40** and peripherals **18**, all being connected to a system bus. The high-level memory module **24** is

an example of another memory module that is accessible by the processor **12**. It usually stores programs and/or data for the processor. It can also be a second level cache memory module supporting off-chip memories, but this is not necessarily so. If a cache miss occurs the data can be fetched from the high-level memory module **24** or from other memory modules.

According to one embodiment of the invention the cache memory is treated as a single unit and can be totally closed or open. According to another embodiment of the invention the cache memory includes multiple portions that managed independently from each other. FIG. **1** illustrates the second embodiment.

The cache memory **20** includes multiple portions **22**, **24**, **26** and **28** can receive power independently from each other. Conveniently each portion includes a single cache way, but this is not necessarily so. A portion can include a group of ways, multiple lines that do not form a way, multiple line segments and the like. Usually power consumption can vary more dynamically when a larger number of segments or lines is involved. Thus, the method can be better suited to cache memory portions that include multiple lines.

Power is provided to power switching circuitry **30** that can, in response to control signals from a power management circuit **40**, supply power to one or more portion of cache memory **20** or power down one or more other portion of cache memory **20** during a low power mode. The power switching circuitry **30** usually includes one or more transistorized switches.

The power switching circuitry **30** is connected to power management circuitry **40** that is adapted to determine whether to power down at least the portion of the cache memory during a low power mode. The determination is responsive to a relationship between: (a) an estimated power gain resulting from powering down the at least portion of the cache memory during the low power mode, and (b) an estimated power loss associated with powering down the at least portion of the cache memory during the low power mode.

The power management circuitry **40** can include hardware components, a combination of software and hardware components and the like. It can cooperate with processor **12** that in turn executes a set of instructions that assist in the determination. The assistance can include estimating either one of the power loss or the power gain.

According to one embodiment of the invention the power management circuitry **40** includes a counter **29** that counts the amount (DI) of dirty information by counting the dirty bits or otherwise estimated the amount dirty data stored within the cache. The estimation can be based upon a comparison between the amount (W) of write operations to the cache by the processor and the amount of write back operations (WB) from the cache to the high level memory. This calculation also takes into account the amount of information written during a write back and write operation. For simplicity of explanation it is assumed that the amount of information that is written during each operation is the same.

DI reflects the amount of valid dirty information. When calculating DI, and especially when calculating the amount of dirty bits, the power management circuitry **40** is also responsive to the valid bits associated with the dirty bits. According to another embodiment of the invention a dirty bit is reset once the information is not valid.

DI=W−WB, and the power management circuitry compares DI to a power gating threshold TH. If DI<TH then the cache memory is flushed and is powered down during a low power threshold.

According to various embodiments of the invention TH (and DI) can be calculated per cache memory portion. TH can be updated in response to various parameters such as temperature, clock frequency, cache statistics, and device behavior.

TH can stored within the power management circuitry **40** but this is not necessarily so.

FIG. **2** illustrates a cache memory **20**, according to an embodiment of the invention.

Cache memory **20** includes multiple ways **22-28** and a cache memory control logic **21**. Each way can powered independently from the other ways.

Assuming that a write back policy is applied, the cache memory control logic **21** updates dirty bits (collectively denoted **25**) to reflect that information was written to the cache memory. The dirty bits also reflect is that information was written to the high-level memory or not.

Conveniently, the cache memory control logic **21** counts, by counter **29**, the amount of asserted dirty bit. The cache memory control logic can also compare, for example by comparator **27**, between the amount of write operations to different locations in the memory and between the amount of write back operations from different locations of the cache memory. This, if a certain cache line (or segment) was involved in different write operations the counter will count only one of these different write operations.

The counter is conveniently reset whenever the cache memory **20**, or the relevant portion is powered down.

FIG. **3** illustrates method **100** for power reduction according to an embodiment of the invention.

Method **100** starts by stage **110** of receiving a request to enter a low power mode. The request can be initiated from various components of a device. Typically, a device includes either a single device power management entity or multiple entities that can determine, for the whole device or for parts of the device, to enter into a low power mode.

Stage **110** is followed by stage **140** of determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain resulting from powering down the at least portion of the component during the low power mode period and an estimated a power loss resulting from powering down the at least one portion during the low power mode period.

For simplicity of explanation stage **140** is split to three stages: stage **120** of estimating a power gain resulting from powering down the at least portion of the component during the low power mode, stage **130** of estimating a power loss resulting from powering down the at least portion of the component during the low power mode, and stage **145** of determining whether to power down the at least portion of the component in response to a relationship between the power gain and power loss.

Stage **120** includes estimating a power gain resulting from powering down the at least portion of the component during the low power mode. It is noted that this estimation can be done each time a request for entering power down is received but this is not necessarily so.

According to various embodiments of the invention this estimation is made in a periodical manner, in a random manner, in a pseudo-random manner, and/or in response to the occurrence of events. These events can include a certain battery level, a change in a clock frequency supplied to the component, and the like.

According to an embodiment of the invention the power gain is predefined and can be provided by the manufacturer of the device that includes the cache, can be downloaded over a network, and the like.

According to yet another embodiment of the invention the estimation is responsive to the duration or other characteristics of the low power mode. This duration can be estimated by monitoring the operation of the device and trying to estimate the duration of the power down period. This monitoring can be performed over long periods (hours, days, weeks and even more), or over shorter periods.

Stage **170** illustrates this monitoring stage. Stage **170** can be executed in parallel to other stages of method **100**, but can also be executed serially to these other stages. Typically, due to the complexity and length of stage **170** it is implemented by software or a combination of software and hardware.

The estimation can be highly effective if the low power period follows a repetitive (or otherwise significantly predictive) pattern. For example, short low power periods of about the same length can occur between pressing of a mobile phone keys, when dialing a number, writing SMS messages or otherwise contacting the mobile keys. When, for example, the user dials a number then the amount of dialed digits and hence the amount of short low power periods can be estimated. Longer low power modes can occur between one phone conversation and another.

Conveniently, the component is a cache memory. Conveniently, the cache memory includes independently powered portions. There portions can be cache lines, cache ways, group of lines that differ from cache ways, groups of ways and the like.

Stage **120** is followed by stage **130** of estimating a power loss associated with powering down the at least portion of the component during the low power mode.

This power loss indicates the power consumed during information retrieval operation from the cache memory (for example during a flushing operation) and information retrieval to the cache memory.

Conveniently the estimation is responsive to the amount of dirty information that is stored in the cache. Dirty information is information that is stored in the cache but not in another, usually higher level, memory unit. The presence of such information is indicated by a dirty bit. Said bit is also referred to as sticky bit.

According to an embodiment of the invention an amount of dirty lines is compared to a threshold representative of the estimated power gain in order to determine the relationship between the power loss and power gain.

Stage **130** is followed by stage **145** of determining whether to power down the at least portion of the component in response to a relationship between the power gain and power loss. Conveniently, if the power gain is greater than the power loss the cache memory is shut down during the low power mode. If the gain equals the loss then the component can be shut down or remain powered up.

Stage **140** is followed by stage **150** of selectively providing power to at least a portion of a component of an integrated circuit during a low power mode.

According to an embodiment of the invention the cache memory includes multiple independently powered portions. It is assumed that these portions are ways. In such a case stages **120** and **130** can be applied for each portion to determine whether or not to power down that certain components. Conveniently, certain portions are allocated for storing dirty information and the estimation is applied on these portions alone. For example, assuming that four thresholds (usually the same threshold) TH**1**-TH**4** are associated with portions **22-28** of cache memory **28** and that the amount of dirty bits within each portion is denoted DI**1**-DI**4** respectively then the method include performing four comparisons between TH**1**-

TH4 and DI1-DI4 are for each portion, determining whether to perform a power down during the low power mode.

According to an embodiment of the invention method **100** also includes stage **160** of generating cache statistics. These statistics can affect the estimation of power loss. For example, if after a low power period during which the cache memory was not powered down, the miss rate is high then the method can decide that it should power down the cache memory during the next power period.

According to another embodiment of the invention the power loss and/or the power gain can be evaluated in response to various parameters that affect the power consumption rate. These parameters can include temperature, the frequency of the clock signal provided to the component and the like. Higher temperatures usually result in lower clock frequencies and higher energy consumption. Clock frequency can be changed when applying power reduction techniques such as DVLS. The frequency of clock signal that is provided to the cache affects its power consumption.

Variations, modifications, and other implementations of what is described herein will occur to those of ordinary skill in the art without departing from the spirit and the scope of the invention as claimed. Accordingly, the invention is to be defined not by the preceding illustrative description but instead by the spirit and scope of the following claims.

We claim:

1. A method for power reduction, the method comprises:

selectively providing power to at least a portion of a component of an integrated circuit during a low power mode; and

determining whether to power down the at least portion of the component in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode.

2. The method according to claim **1** wherein the component is a cache memory.

3. The method according to claim **2** wherein the component is a cache memory that comprises independently powered portions.

4. The method according to claim **2**, wherein the determining is responsive to an amount of information that is stored only in the cache memory.

5. The method according to claim **1**, wherein the determining is responsive to an amount of information associated with dirty bits.

6. The method according to claim **1**, wherein the determining of power gain is responsive to an amount of information stored only within a cache memory and the location of said information within various independently powered portions of the cache memory.

7. The method according to claim **1**, generating cache statistics, and wherein the determining is responsive to the generated cache statistics.

8. The method according to claim **1** wherein the determining is responsive to power consumed during information retrieval.

9. The method according to claim **1** monitoring integrated circuit behavior, and wherein the stage of determining is responsive to results of the monitoring.

10. The method according to claim **1** wherein the stage of determining is responsive to temperature.

11. The method according to claim **1** wherein the stage determining is responsive to clock frequency.

12. A device having power reduction capabilities, the device comprises:

power switching circuitry adapted to selectively provide power to at least a portion of a component of the device during a low power mode, having power management circuitry adapted to determine whether to power down at least the portion of the component during a low power mode in response to a relationship between an estimated power gain and an estimated power loss resulting from powering down the at least portion of the component during the low power mode.

13. The device according to claim **12** wherein the component is a cache memory.

14. The device according to claim **13** wherein the component is a cache memory that comprises independently powered portions.

15. The device according to claim **13**, wherein the determination is responsive to an amount of information that is stored only in the cache memory.

16. The device according to claim **12**, wherein the determination is responsive to an amount of information associated with dirty bits.

17. The device according to claim **12** wherein the estimation of power gain is responsive to an amount of information stored only within a cache memory and the location of said information within various independently powered portions of the cache memory.

18. The device according to claim **12**, further adapted to generate cache statistics, and wherein the determination is responsive to the generated cache statistics.

19. The device according to claim **12** wherein the determination is responsive to power consumed during information retrieval.

20. The device according to claim **12** further adapted to monitor integrated circuit behavior, and wherein the determination is responsive to results of the monitoring.

21. The device according to claim **12** further adapted to receive a temperature indication and wherein the determination is responsive to the temperature indication.

22. The device according to claim **12** further adapted to receive a clock signal frequency indication and wherein the determination is responsive to the clock signal frequency indication.

23. A method for power reduction, the method comprises:

selectively powering down a portion of a component of an integrated circuit in response to the portion having a plurality of dirty bits greater than a threshold number of dirty bits.

* * * * *

FORM 19. Certificate of Compliance with Type-Volume Limitations **Form 19**
July 2020

UNITED STATES COURT OF APPEALS
FOR THE FEDERAL CIRCUIT

CERTIFICATE OF COMPLIANCE WITH TYPE-VOLUME LIMITATIONS

Case Number: 23-1266

Short Case Caption: VLSI Technology LLC v. Intel Corporation

Instructions: When computing a word, line, or page count, you may exclude any items listed as exempted under Fed. R. App. P. 5(c), Fed. R. App. P. 21(d), Fed. R. App. P. 27(d)(2), Fed. R. App. P. 32(f), or Fed. Cir. R. 32(b)(2).

The foregoing filing complies with the relevant type-volume limitation of the Federal Rules of Appellate Procedure and Federal Circuit Rules because it meets one of the following:

☑ the filing has been prepared using a proportionally-spaced typeface and includes 10,187 words.

☐ the filing has been prepared using a monospaced typeface and includes _____ lines of text.

☐ the filing contains _____ pages / _____ words / _____ lines of text, which does not exceed the maximum authorized by this court's order (ECF No. _____).

Date: 04/21/2023

Signature: /s/ Kamran Vakili

Name: /s/ Kamran Vakili

Save for Filing